ISP

Large-Scale Systems Directions: Mid-Year Update - 1984

INPUT®

U-CLS2
1984 c.1

LARGE-SCALE SYSTEMS DIRECTIONS:
MIDYEAR UPDATE - 1984

LARGE-SCALE SYSTEMS DIRECTIONS:

MIDYEAR UPDATE - 1984

CONTENTS

LARGE-SCALE SYSTEMS DIRECTIONS:
MIDYEAR UPDATE - 1984

EXHIBITS

I INTRODUCTION

- INPUT has been forecasting detailed residual values for IBM and software-compatible mainframes since 1977 and for selected peripheral products since 1979. The emphasis of the Residual Value Forecasting Series was always upon analysis and anticipation of significant product development and pricing strategies, rather than upon the mere reporting of used-market prices.

- The inauguration of the Large-Scale Systems Directions program reflects this emphasis. This particular report, Large-Scale Systems Directions: Midyear Update - 1984 ties back to analyses that began in the last report in the Residual Value Forecasting Series and concludes a general overview of large-scale systems directions.

  - Residual Value Forecasts for Large-Scale Systems, a December 1983 report, contained a general analysis of IBM's approach to distributed processing in view of IBM's recently announced 3270 PC, XT/370, and 8150.

  - The first Large-Scale Systems Direction report (Large-Scale Systems Directions: Disk, Tapes, and Printer Systems, March 1984) contained a comprehensive systems view of directions in large-scale peripheral systems.

  - This report contains a general analysis of the shifting role of large-scale systems, and the anticipated changes that can be expected during the 1990s.

- The three reports should be reviewed as a framework within which to view more detailed analyses within specific areas. These analyses will appear in future issues of Large-Scale Systems Directions.

- Chapter II of this report contains the analysis of the shifting role of large-scale mainframes with special emphasis upon the key role of IBM systems software.

- Chapter III contains the midyear update of projected residual values of both selected IBM peripheral systems and large-scale IBM and software-compatible mainframes.

    - The recently announced IBM 3280 magnetic tape subsystem is analyzed in terms of its impact (or lack of impact) on 3240 tape systems.

    - The impact of the 308X series is reviewed, and projected residual values of these new systems are included.

II THE SHIFTING ROLE OF LARGE-SCALE SYSTEMS

A. THE CENTRAL ROLE OF IBM SOFTWARE

1. OBSERVABLE TRENDS

- The most obvious observable trend is that the demand for large mainframes seems to be increasing exponentially. Generations of large-scale processors are being announced with increased frequency. The IBM Sierra will reportedly be announced in late 1984, but the follow-on Summit series is already rumored for 1988 announcement. The quest for MIPS seems unending.

- There is an underlying trend that fuels the MIPS requirements--IBM software. This is the year of MVS/XA, and DB2 will join IMS on large IBM mainframes; the combination of the two should be enough to keep the reported 50-MIPS Sierra quadruple processor (quad) busy. In fact, if competitors are right, and an IBM dual processor only gets 1.8 times the performance of the uniprocessor, and the quad only gets 1.5 times the dual processor, you will only get 33.75 MIPS anyhow. Details are in Exhibit II-1. IBM has stated that MVS/XA will run out of gas in the late 1980s. MVS/XB is being predicted for announcement in 1988--just in time to use up all the MIPS the Summit series can provide.

- IBM software has sold a lot of iron, and it has evolved over the last 20 years based on concepts more appropriate for the bygone era when processing power

EXHIBIT II-1

THE COST OF CONNECTION
(Projected Sierra Processor Power)

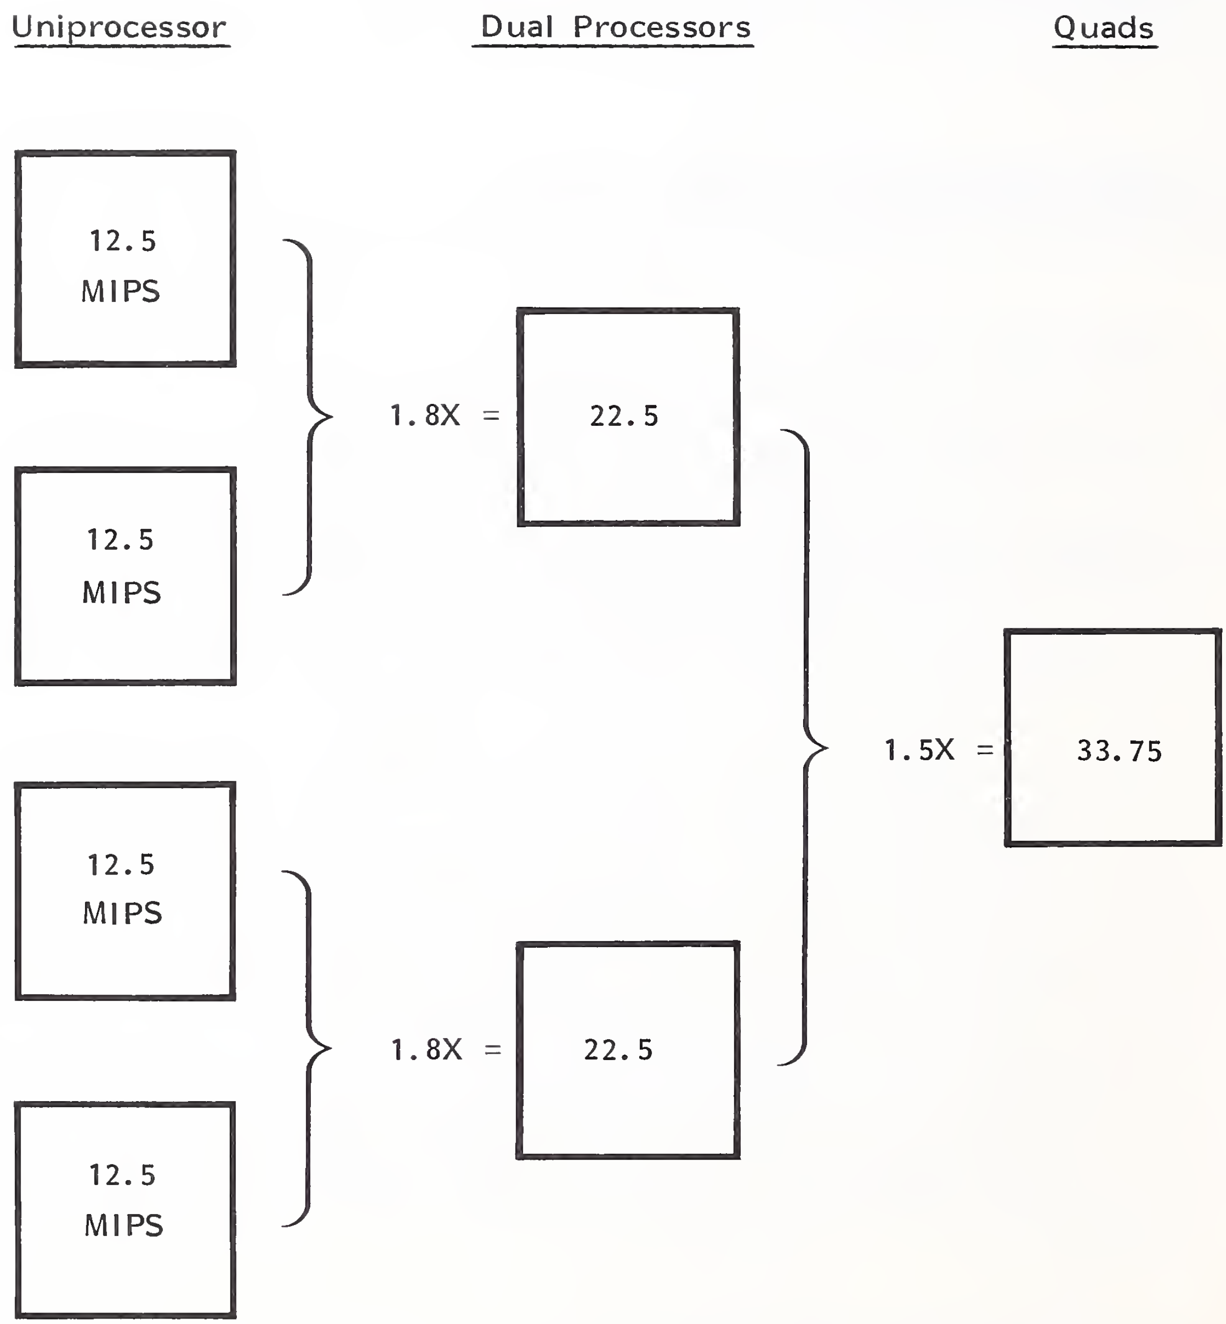

was expensive, random access memory was limited, and jobs were entered through a card reader. The software was not originally designed for today's network environment, and the hardware was not designed to run today's software. As someone familiar with the history of System/360 stated recently: "XA really stands for extended accommodation." That has been the case over the last 20 years as multiprogramming, interactive processing, DBMSs, virtual storage, and multiprocessing capabilities have been added to what essentially started as a stacked job monitor. However, the IBM trend has been consistent in one regard--the progressive centralization of systems on ever-larger mainframes under the control of ever-more-complex systems software.

- Parallel with this IBM emphasis upon centralization has been a technological trend toward distributed processing--first to minicomputers and later to microprocessors. The economies of distributed processing have been clear for well over ten years; eight years ago INPUT recommended the centralization of processing on large central mainframes, to be followed by "the orderly distribution of processing" to minicomputers and intelligent terminals. Despite lip service to distributed processing, IBM's primary strategy has been to preserve (and enhance) the dominant role of large mainframes in its Systems Network Architecture (SNA). (IBM's distributed processing strategy was analyzed in INPUT's Residual Value Forecasts for Large-Scale Systems, December 1983.)

- One additional trend has been to refer to commercial computer applications as data processing systems, management information systems, decision support systems, and expert systems. To date, this trend has represented refinement in terminology more than it has change in substance, but the trend nevertheless gives some indication of what is expected of large mainframes when they are not running payroll and accounts receivable systems.

2. OBVIOUS IMPLICATIONS

- It is obvious that large mainframes are doing something more than the arithmetic required to run business enterprises. A single IBM 3084 can perform more arithemic operations than the entire labor force of the United States (assuming a 40-hour week for the work force and 18 shifts per week for the 3084.) In fact, a single 3084 should be able to perform all the arithmetic operations necessary to process payroll and accounts receivable for the entire United States.

- Even allowing for processing data input and formatting reports, it is obvious that a large mainframe has more than enough processing power for even the largest corporations' required commercial applications. Ask most companies what their major mainframe applications are now and what they were 20 years ago, and you will find that there really hasn't been much change. However, installed processing power in the corporate data center has increased by a factor of at least 100 over the last 20 years.

- This is not quite so difficult to understand if it is recognized that less than 10% of mainframe execution time is spent running user-written code. The rest of the time the processor is busy driving the operating system or a data base management subsystem. Obviously the operating system must be doing a lot of very important work to justify this expenditure of resources. Indeed, the three broad objectives of operating systems can hardly be disputed. They are:

    - Maximum ease of use.

    - Maximum use of equipment (thereby increasing efficiency and reducing the cost per user by sharing resources).

    - Effective development, testing, and introduction of system function without at the same time interfering with service.

- It is difficult to imagine the terrible state of affairs that would exist today if we did not have MVS/XA to address these objectives. By implication, systems would be much more difficult to use, less efficient in their use of hardware, and more disruptive whenever changes were introduced.

- Operating systems are concerned with some very practical problems that have presented a considerable intellectual challenge over the past 20 years. A great deal has been accomplished toward solving these problems. Since such a substantial portion of total computer resources is devoted to the solutions, it is important to understand roughly what they are:

  - The first is referred to as process; it concerns the problems of concurrency (both device and programs) in a multiprogramming environment. The timing required to permit user programs to execute in parallel without interfering with each other is extremely intricate when resources must be shared.

  - Memory management in a multiuser environment creates an extremely complex problem, especially if it is deemed necessary to create the illusion that enormous amounts of main memory exist for each user. Virtual storage concepts have been developed as the primary means of solving the memory management problem.

  - The protection and security of shared data has created problems of access control and, more recently, information flow. Underlying these technical problems is the potentially more important question of privacy.

  - Scheduling and resource management of large systems in order to achieve the efficient use of hardware resources has been an especially irritating problem for computer scientists. This is true because it has been found that techniques from operations research and industrial engineering (such as queuing networks) sometimes provide answers that

conflict with architectural concepts of both hardware and software. (Hardware vendors seem content to leave the more complex scheduling and resource management problems to computer scientists and users. This is not surprising.)

- The combination of all of the above presents a specific problem of how operating systems themselves should be structured (system structure). IBM's highly centralized approach has seen the emergence of virtual machine (VM) concepts. Computer scientists are quick to point out that the VM does not introduce a new concept of systems design, but the same scientists readily admit that the concept of VMs facilitates the introduction of change.

- Considering the changes in hardware technology that have occurred over the past 20 years, it would seem obvious that many of the functions previously assigned to software could perhaps be more easily and economically assigned to hardware. Indeed, it has been necessary to provide hardware to implement certain operating systems functions (such as virtual storage).

- In addition, the basic premise of the shared use of expensive computer power must be questioned since it has become more economical to provide each individual with enough personal computer power to satisfy a high percentage of processing requirements. In other words, can many of the problems solved by operating systems be avoided? This is an especially pertinent question as regards the process, and resource and scheduling problems mentioned above.

- The conclusion must be reached that large mainframes (and associated software) are required primarily to maintain large data bases and to provide processing power for calculations that cannot be realistically performed on distributed systems (minicomputers and microprocessors). For example, calculations that would require hours or days on a personal computer would not be deemed realistic. The question now becomes whether an IBM 3084 operating under MVS/XA is an effective or even appropriate way to provide these services.

3. OBSCURE EFFECTS

- In the early days of OS/360 a system programmer observed: "When you consider what it is doing, it is amazing it works at all." Since then IBM operating systems have become so complex that no individual can comprehend what they do, much less how they do it. There has been general consensus that IBM operating systems do an awful lot but that they are not very efficient at doing what they do. The large overhead of IBM operating systems has come to be an accepted fact of life. This blind acceptance is dangerous because it distorts reality; it is necessary to question what software does to you as well as what it does for you.

- The case of virtual storage is a classic example of obscure effects that continue to this day. Since virtual storge remains central to IBM's (and the industry's) approaches to memory management, it is important to understand both its history and current status.

  - IBM's research on virtual storage started when the largest memory on IBM systems was approximately a quarter of a megabyte. The discovery of "locality of reference" in both matrices and the inner loops of FORTRAN programs convinced IBM that paging was more efficient than other methods of overlaying.

  - Questions concerning locality of reference in commercial applications (where sorting and decision tables (or structures) are common) were not considered, understood, or researched.

  - The eventual problems with "thrashing" (described as a nearly total collapse of processing efficiency) in early implementations of virtual storage systems were described as "a counter-intuitive mystery," but the impact was quite visible; that is, processing of user programs virtually stopped as the system swapped programs and data in and out of main storage.

- While performance degradation is less visible today, it can be assumed that the effect on throughput remains. There is an intuitive attraction to getting work into the system, processing it, and getting it out. The advisability of pushing short jobs through the system as rapidly as possible was proven by research in the late 1960s. Virtual storage systems assure that the process of bringing programs and data together for processing will be an extraordinary exercise in control and scheduling.

- Today, main storage has become large enough and cheap enough to meet the needs of practically all users, and the overlay problem is no longer sufficient reason to justify the use of virtual storage. The justification for virtual storage is now stated to be the built-in protection and process isolation mechanisms that are inherent in systems such as MVS.

- XA is required in order to run MVS. Operating systems have taken on a life of their own. It can truly be stated that large IBM mainframes exist to run IBM operating systems, and operating sytems such as MVS/XA exist only to run large mainframes. The current situation can only be depicted by an M. C. Escher-type drawing in which the user climbs a circular pathway that leads both up and down.

- IBM operating systems are an enigma, but it is important to understand their value now--before IBM turns its full attention to the protection and security of operating systems. There are disturbing signs that users (and systems houses) are going to be further removed from ever being able to understand or improve the performance of systems software.

  - IBM's current trend toward withholding source code is certainly not designed to facilitate understanding or performance improvement (regardless of motivation).

- The trend to preconfigured operating systems (generated by IBM) can also be considered a mixed blessing. While eliminating the need for systems programmers may be hailed as a means of increasing productivity, it also eliminates the primary source of user understanding of IBM hardware/software systems. (Beware whenever IBM states that "the user need not concern himself" or "it will be transparent to the end user," since history has shown that this is precisely when users should concern themselves.)

- Once IBM really addresses the protection and security problem, the mystery of what the operating system does and how it works will be sealed forever. The rule will be: "You can't expect your system to be secure if you know what we are doing to make it secure for you." Even curiosity about what is going on will be suspect, and the mystery will become sacred.

B. LARGE HOST SYSTEMS

I. ENORMOUS DATA BASE MACHINES?

- IBM corporate strategy is essentially large mainframe oriented. Minicomputers have never assumed their proper place under SNA, and micro processors (personal computers) are to be viewed as intelligent workstations heavily dependent upon large host systems. In fact, IBM is counting on intelligent workstations to control the demand for minicomputers and to rejuvenate the demand for large mainframes. (See Residual Value Forecasts for Large-Scale Systems, December 1983.)

- The primary role for large central mainframes can best be described as enormous data base machines. The projected structure of distributed data

bases was presented in INPUT's Large-Scale Systems Directions: Disk, Tape, and Printer Systems, March 1984. There seems to be little question that IBM is confident that intelligent workstations will make enormous demands upon central host systems for data. In other words, IBM has already prescribed a new role for large-scale systems.

- This section will examine this large-scale systems role from an IBM strategic point of view, and it will examine the effects that can be expected as these large and complex systems continue to grow. This section will analyze the question of whether such systems are economically justified or even technically feasible.

2. MAINFRAME MIPS AND STORAGE BITS

- For 1983, IBM reported $10.7 billion in revenue from medium- and large-scale systems, $11.0 billion from standalone peripherals (tapes, printers, and disks primarily), $8.0 billion from small systems/office products, $4.6 billion from maintenance services, $2.3 billion from systems and applications software, and $3.6 billion from "other".

  - Mainframes and standalone peripherals account for a total of $21.7 billion; add appropriate shares of maintenance and software revenues (approximately $4.5 billion), and this total comes to $26.2 billion or 65.2% of IBM's gross revenue of $40.2 billion.

  - The revenues from standalone peripherals, primarily magnetic disk storage, were increasing more rapidly than from any other area (27.3% increase over 1982). It has been INPUT's position for some time that magnetic disk storage represented a key growth area for IBM in the 1980s, and the increase in 1983 revenues confirmed that opinion.

  - Buried somewhere under small systems/office products are personal computers. While financial analysts may feel that PCjr's lack of

market acceptance is a severe blow to IBM's overall strategy, or that the lowering of personal computer prices may seriously threaten IBM's revenues, there is no question about the focus in Armonk, and that the focus is on large host systems with lots of disk storage.

- Both the distribution of data bases and IBM's IMS/DB2 data base architecture were presented in Large-Scale Systems Directions: Disk, Tape, and Printer Systems, the March 1984 report that was previously cited. The functioning of such host data base systems from an operating systems point of view is presented in Exhibit II-2.

  - The projected (IBM) distributed data base environment under the centralized control of large host processors implies a return to a batch environment.

    - The extract facilities of DB2 and the building of relational tables will require batch job scheduling--i.e., some jobs may only run for minutes and others may require hours (or days).

    - While IBM's proposed operating environment does not permit updating of IMS data bases directly from DB2, it is inevitable that changes from DB2 applications will be batched and run against IMS data bases. (In addition, distributed processors will have the effect of batching changes for central data bases under any circumstances.)

    - Mirror data bases and backup is the name of the game in an environment where the sale of disk storage is the key to revenue growth; there will be backups of backups unless careful planning is done.

    - The whole micro-mainframe link concept is to provide for the extraction and downloading of planning and personal data

EXHIBIT II-2

LARGE HOST DATA BASE MACHINES

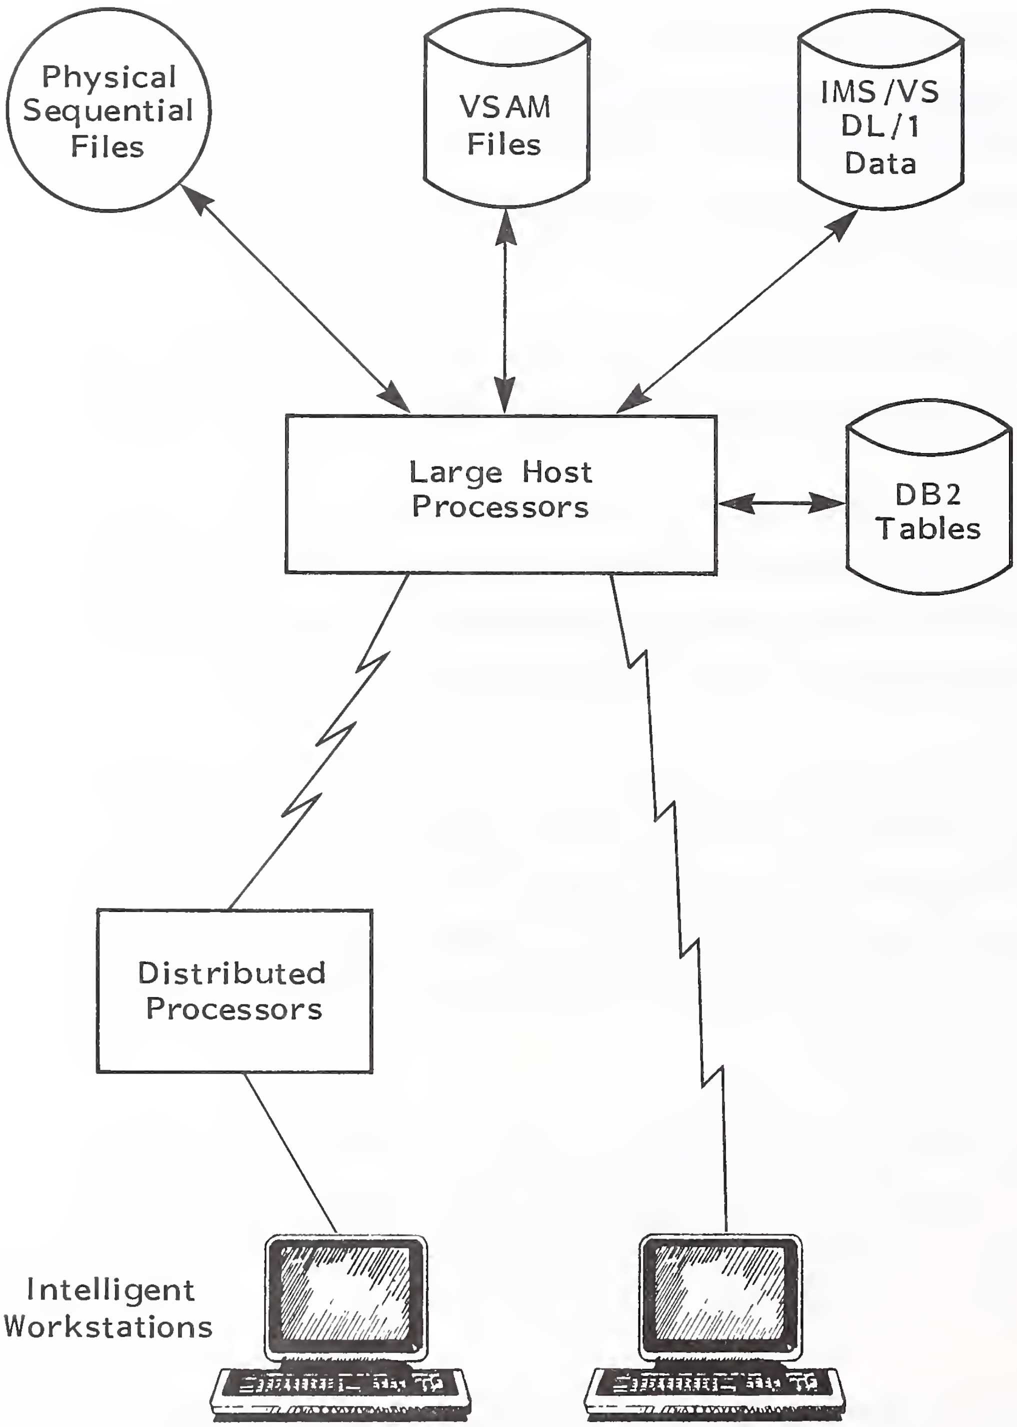

1 Schedule Batch Requests

- (a) Extracts from File to Build DB2 Tables.
- (b) Updating of Data Bases.
- (c) Backup of Host & Distributed Data Bases.
- (d) Extract and Transmission of Data Bases.
- (e) Heavy Computation.

2 Interactive Processing

- (a) IMS Data Bases.
- (b) DB2 Tables.

3 Protection & Security

- (a) File & Data Base Access.
- (b) Data Base Update.
- (c) File Syncronization & Integrity.
- (d) Certified Files & Data Bases.
- (e) Monitor Information Flow.
- (f) Control Encryption.

bases. (If there aren't batch requests on the host, there will be problems that will be described later.)

- Heavy computation implies batch processing (although some users may expect interactive response to requests for heavy computation).

- Against this batch environment it will also be necessary to provide interactive support (or vice versa). For the type of environment envisioned, processing of batch and interactive cannot be conveniently separated (or scheduled separately). An operator at an intelligent workstation may initiate requests for a batch run against a sequential file, an IMS inquiry and JOIN and SELECT commands against relational tables under DB2 and have them running in parallel. Plus, the intelligent workstation doing all the above may be in the president's office.

- The distributed data base environment that is anticipated because of micro-mainframe links is going to bring protection and security to the fore in everybody's mind--and for good reason.

  - The traditional problems of access can only get worse as the new generation of hackers come on-line with super microprocessors to help them. There will soon be as much (or more) concern about internal access as there is now about access from competitors or hobbyists.

  - It is assumed that intelligent workstations will generate increasing amounts of data and information and that there will be a desire to update data bases and files; contamination may become more of a problem than access unless updates are rigidly controlled.

- The well-known concerns about sychronization and integrity will still be justified, even if only authorized persons are permitted to update files; data base management is not going to become any easier.

- Some automatic means of file and data base certification across data bases (IMS-DB2 for example) must be established; this is a nontrivial problem when considered in the projected distributed data base environment.

- The problem of certification is closely associated with the problems of information flow control. These problems currently trouble information scientists and have not been solved. (Simply stated, the question is, How can you control unauthorized information leakage based on authorized use of specific data and imaginative analysis of the data? It is a good question.)

- Encryption hasn't become a problem yet because most users (even those with sensitive data) have not taken the security problems seriously, despite exaggerated and sensational publicity given to "computer crime."

- When IBM has its solution to the security and protection problems incorporated in its large host systems, the whole area is going to receive a lot of high-level attention. Of course, you can be sure the solution is not going to be cheap.

• Even this general description of the large host data base machines should point out quite clearly that IBM's distributed processing strategy is going to generate enormous demands for mainframe MIPS and storage bits. In fact, INPUT has recently forecast that IBM will be able to maintain its traditional growth and become a $100 billion company by 1990 without drastic shifts in emphasis from its current large-scale mainframe, magnetic disk, and intel-

ligent workstation strategy. In other words, SNA (and distributed processing) will continue to evolve at what some consider to be an excruciatingly slow pace.

- INPUT predicted in its last large-scale systems report that IBM will be successful in slowing acceptance of optical disk systems that might effect magnetic storage revenue (or facilitate electronic offices) until the 1988-89 time frame.

- During recent research on another subject, an IBM employee stated: "We have made a lot of money on large mainframes for a long time, and we think it can go on forever." (There are those in IBM who view with alarm potential technological effects on IBM strategy.) INPUT assured the IBM employee that INPUT felt IBM would, in fact, continue to make a lot of money from large-scale systems through the 1980s.

- The fact of the matter is that SNA makes more sense at this point in time than it ever has. The controlled distribution of processing onto minicomputers has been technically possible and economically justified for nearly 15 years, and SNA has fought that prevailing trend for the last 10 years. However, the proliferation of microprocessor-based systems and personal data bases demands central control if chaos is to be avoided. IBM certainly intends to integrate these items under the great SNA umbrella. This integration probably is the only way to go.

- In addition, the operating systems environment described above is the one IBM has always aimed at; that is, IBM assumes an underlying requirement for batch processing and that interactive systems will be run on the same system. While the UNIX buffs are correct in stating that MVS (and its predecessors) were really batch systems with communications and terminals added, it will be interesting to see what happens when and if a system designed for interactive processing is installed in the environment depicted in Exhibit II-2.

- There is only one problem with all this, and that is that the processing burden of IBM operating systems and their data base subsystems may have finally reached the point where processing engines may not be available to drive them. It is certainly not difficult to picture Sierra being overburdened as soon as it comes off the production line.

3. THE LIMITS OF GROWTH

- Earlier this year, INPUT prepared two executive bulletins on the entropy (the natural tendency toward disorder) associated with data and information in the environment described in this report. Our fear was that as central data bases grew in size and complexity, and personal data bases proliferated, chaos would inevitably develop because sufficient energy (human and computer processing power) would not be available to maintain order (or organization). Since these bulletins were published, additional research has been done that supports this intuitive concern:

    - Comprehensive research on IBM software directions has revealed that problems of data entropy are just now being recognized at Yorktown Heights. There is no awareness or understanding of the problem vis-a-vis establishing IBM software directions, and certainly none with regard to implementing current operating or data base systems.

    - The highly centralized nature of IBM's approach to distributed processing (essentially that described by INPUT in this series of reports) will place unanticipated burdens upon host systems as central data bases grow (and become more flexible), and as new functions are added (especially those associated with protection and security).

    - The probability of catastrophic central systems failure exists unless the central data facility is managed with extreme care. This implies central IS control not only of the central data base but also of the demands made upon the system for data. In other words, just because

someone is authorized to have access to data does not mean that all requests can be serviced (even those from the president's office). There is no indication that IBM will provide adequate means of screening such requests.

- Even highly centralized control of corporate data does not ensure the quality of information flow. That will require substantial redundancy in communications, with resultant demands upon the central system.

- In some ways, failure of the central system is preferable to the undetected deterioration in the quality of data and information. It is probable that most corporations would continue to function without the volumes of planning data from the central facility. Most operating decisions are still made based on relatively simple data and analysis, which are available at the local level. Executive decisions remain primarily intuitive in any given instance.

• This brings us to the question of decision support systems and the demands such systems will make on even the largest computer system. The availability of enormous quantities of data--even assuming chaos can be averted--implies that analysis tools that go beyond spreadsheets must be applied to assist the decision maker. The current emphasis upon knowledge bases and expert systems is only the logical extension of current decision support systems. The implementation of expert systems in the general business environment is currently restricted not only by the availability of data and information but also by models that facilitate analysis.

  - Even relatively simple economic models (such as supply and demand) defy mathematical description. John von Neumann expressed the need for a breakthrough in mathematics (<u>Theory of Games and Economic Behavior</u>) over 40 years ago, but progress has been extremely slow.

- In addition, the two areas in which analytical tools and models are being developed--artificial intelligence and operations research--exhibit similar disturbing properties. Specifically:

  - The algorithms of operations research depend heavily upon probability theory and the computational cost increase faster than any finite power of n. For example, 100! comparisons exceeds the capacity of any computing resource on earth.

  - The same type of limitation exists in artificial intelligence where the number of choices increases exponentially (as in playing chess).

  - It is anticipated that relational data bases underlying knowledge bases to feed such models will only compound the problem.

- Thus, it appears that many current "solutions" to improve the decision-making process may exceed not only current systems and fifth-generation systems but also any computing facility that can ever be constructed. These limits of growth in large computer systems should be recognized because current IBM hardware/software systems are pointed in a direction that will reach those limits soon!

C. DISTRIBUTION OF FUNCTION

- Hans J. Bremerman, who derived the physical limits of data processing over 20 years ago, has stated:

  - "We call an algorithm transcomputable if its computational cost exceeds all bounds that govern the physical implementation of algorithms.

- "It can be shown that the exhaustive search algorithm for chess is transcomputable. The same is true for many algorithms of artificial intelligence and operations research. In fact, any algorithm whose computational cost grows exponentially with a size parameter n is transcomputational for all but the first few integers n.

- "This is a rather disturbing thought and many people have chosen to ignore it."

- IBM has not only chosen to ignore these facts but also has proceeded on the assumption that the problem is really to create demand for large-scale processing power as processor technology gets faster and costs get lower. This has been done with complex systems software that manages to consume MIPS faster than IBM releases new technology, and by keeping most processing on the central host. The processors that have been developed are an exaggerated example of overkill in terms of the simple arithmetic required for accounting-type applications, and in terms of the operating systems overhead that has been employed to keep the system busy.

- Now, with the role that has been defined for the large-scale systems of the 1980s--the management of large data bases and the computation to support decision making--the total hardware/software system runs the high risk of failing for lack of central processing power. All that is needed is:

  - The development of a prototype that uses a relational data base system on an intelligent workstation and then goes into operation against corporate data bases (and files) using DB2 on the mainframe.

    - A system that can run a "traveling-salesman-type" problem on a PC for 10 cities and then run the same problem for 100 (or even 50) cities on the 3084.

- The limitations and costs of large-scale IBM- and software-compatible systems will become readily apparent. Both the big general-purpose engines and the supporting software will be replaced during the 1990s, when differentiation and mechanization of current large-scale hardware and software functions occur on a massive basis.

I. PROCESSING

- It has long been known that general-purpose computers of IBM/360-370 architecture, with or without the burden of IBM software, cannot meet the processing requirements of certain potential applications. The classic example is numerical weather prediction, which still requires approximately a 100 to 1000-fold processing increase on even the supercomputers. This ever-growing class of problems has been differentiated from the mainstream and left to more specialized processors from CRAY and CDC.

- It also has become recognized that minicomputers are three to four times more cost-effective than large processors on another set of problems (studies for nuclear reactors, solid-state devices, aircraft design, and petroleum reservoirs) in which turnaround is not a major consideration. Once again a single processor, with minimum software burden and operating on a specific problem, results in a substantially more cost-effective operation.

- More recently, microprocessors have demonstrated that they are substantially more cost-effective for any number of accounting and statistical-type operations (to say nothing of text editing and word processing) that are representative of the commercial environment.

- The problems addressed by large-scale operating systems (process, memory management, scheduling and resource management, general systems structure, and even protection and security) are all alleviated or eliminated once each individual and/or each application is assigned an individual processor. In other words, the assumption of scarce, shared commodities has been replaced by a projection of abundance.

- Looking back at the broad objectives of operating systems, it seems apparent that some change of attitude is required.

  - Maximum ease-of-use can certainly be achieved in an environment where the processing resource is not shared.

  - Maximum use of equipment is based on the scarce commodity and shared use attitude of the past. It has already been pointed out that the distributed processing environment is more cost-effective, and that the attitude that equipment must be kept busy is a thing of the past. A personal computer is more equivalent to a telephone than it is to a 3084, since it is meant for the convenience of the user and not as a piece of production equipment.

  - The effective development and testing of functions without interfering with service is a problem created by the shared facility and software itself. Processors are going to become more like their smaller hand-held brothers, the calculators. As more functions become available or necessary, the new hardware/software replacement is purchased. For those who might say that the incremental addition of functions under operating systems extends the life of the hardware, it can only be pointed out again that large mainframe life cycles are getting shorter, and operating systems are the primary cause.

- The point is that processors are going to be differentiated, based on use and/or application down to a relatively fine level; there will be specialized processors for array processing, pattern requirements, and communications. There is no reason individual processors cannot be assigned for specific batch and interactive tasks. This can be true whether the processors are geographically or architecturally distributed.

- As far as systems structure is concerned, VM leads into this rather conveniently; virtual-real can be substituted over a period of time for machines as well as for storage. The thing that gets eliminated on the real processor is the burden of a full-blown multiprogramming operating system.

2. PERIPHERAL STORAGE MANAGEMENT

- It was stated earlier in this section that centralized storage managment is highly desirable, if not essential, to a distributed processing environment. This did not imply that all physical storage had to be centralized, and most certainly did not mean that data base management systems are necessarily implemented in software under general-purpose operating systems.

- INPUT presented an argument, in Relational Data Base Developments (August 1983), for the desirability (and even necessity) for data base machines and will not repeat the full analysis here except to state that:

  - The IBM/360/370 architecture is not especially well suited for data base management purposes.

  - There is a need for various data models (hierarchical, network, relational), and performance assistance will be necessary if these models are to co-habit (in the network as well as on a single system).

  - DB2 with IMS or large-scale IBM mainframes will bring the performance problem into focus.

  - Data base machines will be necessary.

- This argument for the architectural distribution of mainframe processing represents both the differentiation and mechanization of the data base management functions. A similar case can be made for the geographic distribution of large data bases. (In other words, an argument can be made for breaking them up into manageable pieces.)

- The operation of the JOIN operation in relational data base systems has the disturbing property of having the computational cost (number of comparisons) increase more rapidly than n (the number of elements in the tables). Thus it will be necessary to limit the size of relational data bases.

- The same argument can be applied to any large data base; for example, it is substantially more cost-effective to break down directory assistance for telephone numbers or the Sears credit file on a local basis than it is to have one massive data base. The fact that there are very large central data bases ignores the current economics of computer/communication networks. The cost justification for distributing data bases is rapidly becoming more compelling.

• Questions of privacy and security are also more appropriately and easily managed. Central processing facilities with large data bases on 3380s cannot easily be disconnected from the network, and data bases of varying sensitivity are "on-line" whenever the system is operational. Distributing sensitive data bases solves a lot of problems.

  - The data bases can be completely isolated from access on the general network.

  - The individuals (or department) using the data base can determine the physical security, encryption, and access requirements based on their specific requirements.

  - The big central target for serious or playful hackers is eliminated.

• When considering backup of data bases, large central facilities create the problem of how you backup the backups.

- It is more effective and economical to have comparable nodes backup each other by using a buddy system. (For example, two branches of a bank will each mirror the transactions and account information of the other.)

- Data bases placed as close to the end user as possible are highly preferable in all regards--provided they can be controlled. Such control becomes simpler when emphasis is placed on communications rather than on data processing. The main thing that must be changed is the fortress-like structure of the corporate data base function, which is currently more concept than reality.

3. DOCUMENT PRODUCTION, STORAGE, AND HANDLING

- Surrounding large-scale systems today are the centralized facilities for producing paper documents--miles and miles of paper documents that must be printed, handled (distributed), stored, and/or destroyed. Intelligent workstations on each desk may not eliminate the need for paper, but as processing and data are distributed closer to users, so will document production be distributed. Handling will thus be eased, and the availability of optical memories will permit economical storage--paper use will be reduced.

- It is probable that some large central facilities will remain as paper mills. For example, Reader's Digest may still be mailing out its annual contest to one million households. However, the major trend of the 1990s will be toward paper reduction, and the big central printing facilities will tend to dwindle away.

D. SUMMARY

- Large central mainframes during the 1980s will tend to become enormous data base machines in support of both operational and decision support systems. Applications based on these data bases will tend to become distributed (trans-

action processing, editing of input, report preparation, simple accounting, and statistical arithmetic) because of the price/performance advantages of minicomputers and microprocessors.

- The centralization of control in large host systems (which has been characteristic of SNA and IBM operating systems in general) is highly appropriate at this time because the rush to link intelligent workstations to mainframes has high potential for the deterioration of data and information quality. In fact, even the original batch orientation of IBM operating systems may be an advantage in such an environment.

- However, the performance burden of IBM operating systems (and DBMS subsystems) in the data base environment that is envisioned may exceed the increases anticipated in processor performance. This is especially true because decision support systems will require complex mathematical models if these large central data bases are to be of any practical value in decision support. The commercial environment will suddenly change from one in which simple arithmetic has sufficed to one in which the tools of operations research and artificial intelligence are applied. Many of these tools require computational capabilities that exceed those of today's supercomputers.

- There is the potential for unanticipated failure of such systems; that is, data bases will grow to the point where they cannot be maintained or used for practical purposes, or a "simple request" for processing of a decision-making model may exceed the host systems' capability (or at least what the requestor would be willing to pay).

- Even without catastrophic failure, performance will dictate the differentiation and mechanization of both hardware and operating system functions (the architectural and geographic distribution of processing), and data bases will in turn be distributed. In the 1990s, this will lead to the massive replacement of large-scale systems as we know them today.

- IBM's tradional concern has been that processor price-performance would improve more rapidly than the users' ability to use such power effectively. Systems software has proved more than adequate at absorbing technological advances in processor performance, but IBM has remained reluctant to off-load large mainframes to any significant degree. This may lead to unfortunate consequences for users in the 1980s as the weaknesses of large-scale hardware/software systems are exposed.

- Users are urged to anticipate the performance problems of the projected large-scale systems environment of the late 1980s. Centralized control of data and information must be exercised in the face of increased user involvement in the systems development process, but this very involvement practically assures the performance problems that have been outlined. Therefore, the orderly distribution of processing and data bases (in advance of IBM's schedule) is necessary if unfortunate surprises are to be avoided. This has been a recurrent INPUT theme for eight years, and the reports this year will emphasize what can be done in this regard.

III RESIDUAL VALUE FORECASTS

A. ANNOUNCEMENTS

- When IBM announced the X models of the 308X series in February, INPUT issued an Executive Bulletin that reached the following preliminary conclusions concerning that announcement:

  - IBM could reduce prices on the 308XX by a substantial amount and still achieve traditional profit levels.

  - The relatively modest price-performance improvement meant a substantial profit boost for IBM, and would provide flexibility in Sierra pricing.

  - It seemed doubtful that the 308XX would be field upgradable to Sierra (and that if it were, the cost of the upgrade would affect the price-performance advantages).

  - Announcement and/or delivery schedules of Sierra would be delayed to permit the 308XX profitability advantages to be exploited.

  - Residual values of the 308X series would be impacted, but not as badly as some analysts originally thought. And, since upgrades had been made attractive, it was probable that the used market would recover after the initial confusion.

- INPUT also promised to review those preliminary conclusions in this report. Since then, IBM has announced a performance improvement package for the 308X of up to 6% greater internal throughput for $16,000 (May 1984). Although this has resulted in additional confusion about the purpose of the 308XX announcement (some analysts are now stating that the 308X is a better deal than the 308XX for those interested in imaginative upgrading), our preliminary conclusions still look reasonably good.

    - The 308XX can be reduced in price substantially and still meet traditional IBM profit objectives. This will be done in connection with the SIERRA announcement and will provide considerable flexibility in terms of both pricing and delivery schedules. (At the same time, those upgraded 308X systems will probably lose a lot of their appeal.)

    - It still seems doubtful that the 308XX will be field upgradable without substantial impact on the improved price-performance of the Sierra.

    - Practically everyone predicts Sierra will be announced in the fourth quarter of 1984, but INPUT does not feel IBM is under any great pressure to announce an entirely new line of large-scale mainframes. INPUT looks for something different to be announced:

        - Perhaps a high-end system for those who find that the performance of the 3084 is not adequate for their needs.

        - Perhaps something to fill the open slot on the 308XX--for example, IBM's version of a data base processor.

        - Under any circumstances, the 308XX is going to be around in some form even after the SIERRA (or TROUT or whatever it is currently being called) announcement.

- Residual values have been affected, but they still fall within the range of the charts published in INPUT's Residual Value Forecasts for Large-Scale Systems (December 1983). (Revised tables of residual values for the 308X series are included later in this section.)

- In April, National Advanced Systems (NAS) announced its AS/80X3 series of mainframes that will replace and extend the former AS/80X0 series. The old series is field upgradable, and NAS hailed this announcement as upholding "NAS's tradition of protecting its customers' investments." Indeed, there is truth in this statement, and residual value forecasts for NAS processors have been adjusted accordingly. (The AS/80X3 series covers a range extending from the IBM 4381-2 through the IBM 3081KX.)

- In June, TRILOGY (Amdahl & Son) announced discontinuance of its attempt to enter the very large-scale, software-compatible mainframe market with a system based on wafer-scale technology. The announcement followed numerous delays in shipment date, and approximately $250 million in financing. It is indeed unfortunate because now there will not be any forced price-performance adjustments on large-scale IBM systems comparable to those resulting from Gene Amdahl's earlier efforts.

- In late March, IBM announced its long-awaited new tape subsystem. While it represented a substantial change in magnetic tape technology, it fell short of INPUT's expectations.

    - Essential features of the IBM 3480 Tape Subsystem are as follows:

        - Three-megabyte-per-second data rate.

        - Linear recording density of approximately 38,000 bits per inch, on half-inch chromium dioxide tape housed in a protective plastic cartridge.

- 18-track thin-film read/write heads.

- Reliability through advanced component technology, new error correction code, and separate microprocessors in each drive and control unit.

- The subsystem attaches to the block multiplexor of the IBM 303X, 308X, 4341, and 4381 processor under MVS/370 and MVS/XA.

- The units consume only half the space and 60% less power and cooling capacity than the 3420.

- General availability will be in the first quarter of 1985.

- It was generally felt that because of cost (a "typical" 3480 subsystem consisting of one controller and eight drives costs $238,000) and conversion problems, the 3480 will not replace the venerable 3420 tape technology.

- The announcement leaves a substantial problem of 3380 disk file dumping and backup for large mainframe sites. This problem will be compounded as more on-line storage is added. There are already rumors of new technology (perhaps optical disk?) to address the problem.

B. USED-MARKET ACTIVITY

- Used-market activity and prices are a major determining factor in the residual value of installed equipment. Exhibit III-1 and III-2 present used market average retail prices for selected IBM peripheral and mainframes (as a percent of IBM list price).

EXHIBIT III-1

USED MARKET AVERAGE RETAIL PRICES FOR SELECTED IBM PERIPHERALS

(As a Percent of IBM List)

Model	1982	1983				1984	Market
	December	March	June	September	December	March	Trend
3330-001	3%	3%	3%	2%	1%	1%	Down
3330-011	4	3	3	2	1	1	Down
3350-A02	52	52	55	38	25	32	Up
3350-B02	53	53	55	38	25	32	Up
3380-AA4	103	101	101	93	93	90	Stable
3380-B04	103	101	101	93	93	90	Stable
3420-003	8	8	8	7	5	5	Down
3420-005	10	10	10	12	12	12	Stable
3420-007	13	17	17	21	28	28	Stable
3420-004	55	55	57	68	63	60	Stable
3420-006	54	50	58	75	70	70	Stable
3420-008	69	67	71	93	93	88	Stable
3480-B22	-	-	-	-	-	112	Down
1403-N01	7	5	4	3	3	3	Down
3211-001	50	55	55	57	52	48	Down
3800-001	63	63	60	58	58	56	Down

The values shown are used-market retail prices. At any given time, three price levels exist:

Retail Price - The amount an end user would pay for the equipment.

Dealer Price - The amount a dealer would pay another dealer to acquire equipment to complete a contracted sales obligation.

Wholesale Price - The amount a dealer would pay to acquire equipment for resale.

The dollar spread between levels is a function of the total value of the transaction. For large processors the wholesale price will typically be 80% to 95% and for peripheral equipment 70% to 90% of the retail price.

EXHIBIT III-2

USED MARKET AVERAGE RETAIL PRICES FOR SELECTED IBM LARGE MAINFRAMES (As a Percent of IBM List)

Model	1982	1983				1984	Market
	December	March	June	September	December	March	Trend
4331-1	58	55	51	42	41	40	Stable
4331-2	63	60	60	60	64	62	Stable
4341-1	78	71	71	67	57	41	Down
4341-2	79	75	75	67	58	55	Down
3031-6	6	5	5	4	3	2	Down
3032-8	7	5	4	4	3	-	Down
3033-N	28	26	24	15	13	6	Down
3033-U	30	28	25	20	14	11	Down
3083-E	-	-	-	-	88	81	Down
3083-B	-	-	-	-	90	82	Down
3083-J	-	-	-	-	90	82	Down
3081-D	90	85	88	85	80	80	Down
3081-K	94	92	92	90	85	85	Down

The values shown are used-market retail prices. At any given time, three price levels exist:

Retail Price - The amount an end user would pay for the equipment.

Dealer Price - The amount a dealer would pay another dealer to acquire equipment to complete a contracted sales obligation.

Wholesale Price - The amount a dealer would pay to acquire equipment for resale.

The dollar spread between levels is a function of the total value of the transaction. For large processors the wholesale price will typically be 80% to 95% and for peripheral equipment 70% to 90% of the retail price.

- The used market for peripherals has not changed substantially from that described in Large-Scale Systems Directions: Disk, Tape, and Printer Systems (March 1984).

  - The 3480-B22 magnetic tape subsystem is currently selling at an early shipment premium awaiting volume shipment in the first quarter of 1985. The market for 3420 models has generally remained stable since the announcement of the 3480.

  - The 3350 is going through another temporary resurgence as the oversupply created by 3380 replacement of installed 3350s has slackened.

- The announcement of the 308XX series of processors has had a generally depressing effect on used-market prices for 308X equipment, but the general used-market trend downwards had already been established last year. The effect of the performance improvement package (coupled with the cheap upgradability that was mentioned earlier) should have the effect of firming up the used market for 308X processor, but this remains a year of general confusion in the large-mainframe market.

C. PROJECTED RESIDUAL VALUES

- Exhibit III-3 and III-4 present the projected used-market retail value in dollars, and the projected residual value of IBM and software-compatible mainframes as a percent of vendor list price. Exhibit III-5 through III-17 graph the range of anticipated values (as a percent of list price) for 1985 through 1989 for selected processors:

  - IBM 4341-2, 4361, 4381, 3083EX, 3083JX, 3081GX, 3081KX, and 3084QX.

EXHIBIT III-3

PROJECTED USED-MARKET RETAIL VALUE AT JANUARY 1

VENDOR	PROCESSOR MODEL	CURRENT LIST 6/1/84	1985	1986	1987	1988	1989
IBM	4331-J01	60920	21322	15839	9747	4264	2437
	4331-K02	90000	57600	43200	29700	10800	6300
	4341-K01	184500	77490	51660	25830	9225	5535
	4341-L02	312000	149760	109200	62400	24960	15600
	4361-0K4	150000	130500	114000	90000	60000	19500
	4381-0K5	200000	180000	164000	140000	96000	40000
	3033-N08	1474000	88440	44220	14740	14740	0
	3033-U12	1964000	157120	157120	157120	157120	157120
	3083-E08	1190000	797300	559300	297500	107100	35700
	3083-B08	1805000	1227400	902500	505400	198550	90250
	3083-J08	2330000	1631000	1188300	699000	326200	163100
	3083-EX8	1190000	1011500	737800	571200	190400	95200
	3083-BX8	1805000	1534250	1173250	902500	342950	180500
	3083-JX8	2330000	1980500	1561100	1188300	489300	256300
	3081-G24	3325000	2161250	1596000	897750	299250	99750
	3081-K24	3855000	2698500	2120250	1349250	539700	231300
	3084-Q32	*					
	3081-GX24	3325000	2859500	1762250	1064000	399000	166250
	3081-KX24	3855000	3469500	2274450	1464900	616800	308400
	3084-QX32	*					
AMDAHL	470-V7		0	0	0	0	0
	470-V8		0	0	0	0	0
	5840-16	2000000	1800000	1040000	700000	340000	180000
	5850-24	2660000	2394000	1436400	984200	532000	266000
	5860-24	2910000	1949700	1455000	989400	436500	232800
	5867-24	3560000	3382000	2171600	1352800	605200	391600
	5868-32	4070000	3866500	2564100	1668700	895400	569800
	5870-32	4520000	4068000	2712000	1762800	858800	542400
	5880-48	5340000	4272000	2776800	1975800	961200	534000
NAS	AS/6620	255000	127500	71400	35700	12750	7650
	AS/6630	341500	191240	112695	61470	27320	17075
	AS/6650	417500	250500	146125	91850	50100	37575
	AS/8023	639000	607050	332280	210870	76680	38340
	AS/8043	1255000	1029100	690250	439250	188250	100400
	AS/8053	1758000	1459140	1002060	632880	263700	158220
	AS/8063	2251000	2251000	1350600	877890	382670	247610
	AS/8083	3506000	n/a	2454200	1507580	701200	490840
	AS/9040	1758000	632880	351600	158220	87900	35160
	AS/9050	2256000	902400	564000	270720	157920	90240
	AS/9060	2729000	1173470	818700	491220	272900	163740
	AS/9070	3706000	1853000	1556520	889440	555900	296480
	AS/9080	4722000	2833200	2361000	1416600	849960	472200

EXHIBIT III-4

PROJECTED RESIDUAL VALUES FOR IBM AND SOFTWARE-COMPATIBLE MAINFRAMES

VENDOR	PROCESSOR MODEL	PROJECTED RESIDUAL VALUE AS PERCENT OF VENDOR LIST PRICE AS OF JANUARY 1				
		1985	1986	1987	1988	1989
IBM	4331-1	35	26	16	7	4
	4331-2	64	48	33	12	7
	4341-1	42	28	14	5	3
	4341-2	48	35	20	8	5
	4361-1	86	74	56	35	10
	4361-2	87	76	60	40	13
	4381-1	88	80	65	42	14
	4381-2	90	83	70	48	20
	3033-N	6	3	1	1	-
	3033-U	8	5	3	2	1
	3083-E	67	47	25	9	3
	3083-B	68	50	28	11	5
	3083-J	70	51	30	14	7
	3083-EX	85	62	48	16	8
	3083-BX	85	65	50	19	10
	3083-JX	85	67	51	21	11
	3081-G	65	48	27	9	3
	3081-K	70	55	35	14	6
	3084-Q	78	61	40	18	9
	3081-GX	86	53	32	12	5
	3081-KX	90	59	38	16	8
	3084-QX	90	65	44	21	11
AMDAHL	470-V/7	7	4	2	1	1
	470-V/8	10	7	5	2	1
	5840	90	52	35	17	9
	5850	90	54	37	20	10
	5860	67	50	34	15	8
	5867	95	61	38	17	11
	5868	95	63	41	22	14
	5870	90	60	39	19	12
	5880	80	52	37	18	10
NAS	AS/6620	50	28	14	5	3
	AS/6630	56	33	18	8	5
	AS/6650	60	35	22	12	9
	AS/8023	95	52	33	12	6
	AS/8043	82	55	35	15	8
	AS/8053	83	57	36	15	9
	AS/8063	100	60	39	17	11
	AS/8083	-	70	43	20	14
	AS/9040	36	20	9	5	2
	AS/9050	40	25	12	7	4
	AS/9060	43	30	18	10	6
	AS/9070	50	42	24	15	8
	AS/9080	60	50	30	18	10

EXHIBIT III-5

RESIDUAL VALUE FORECAST FOR
IBM 4341-2 PROCESSOR

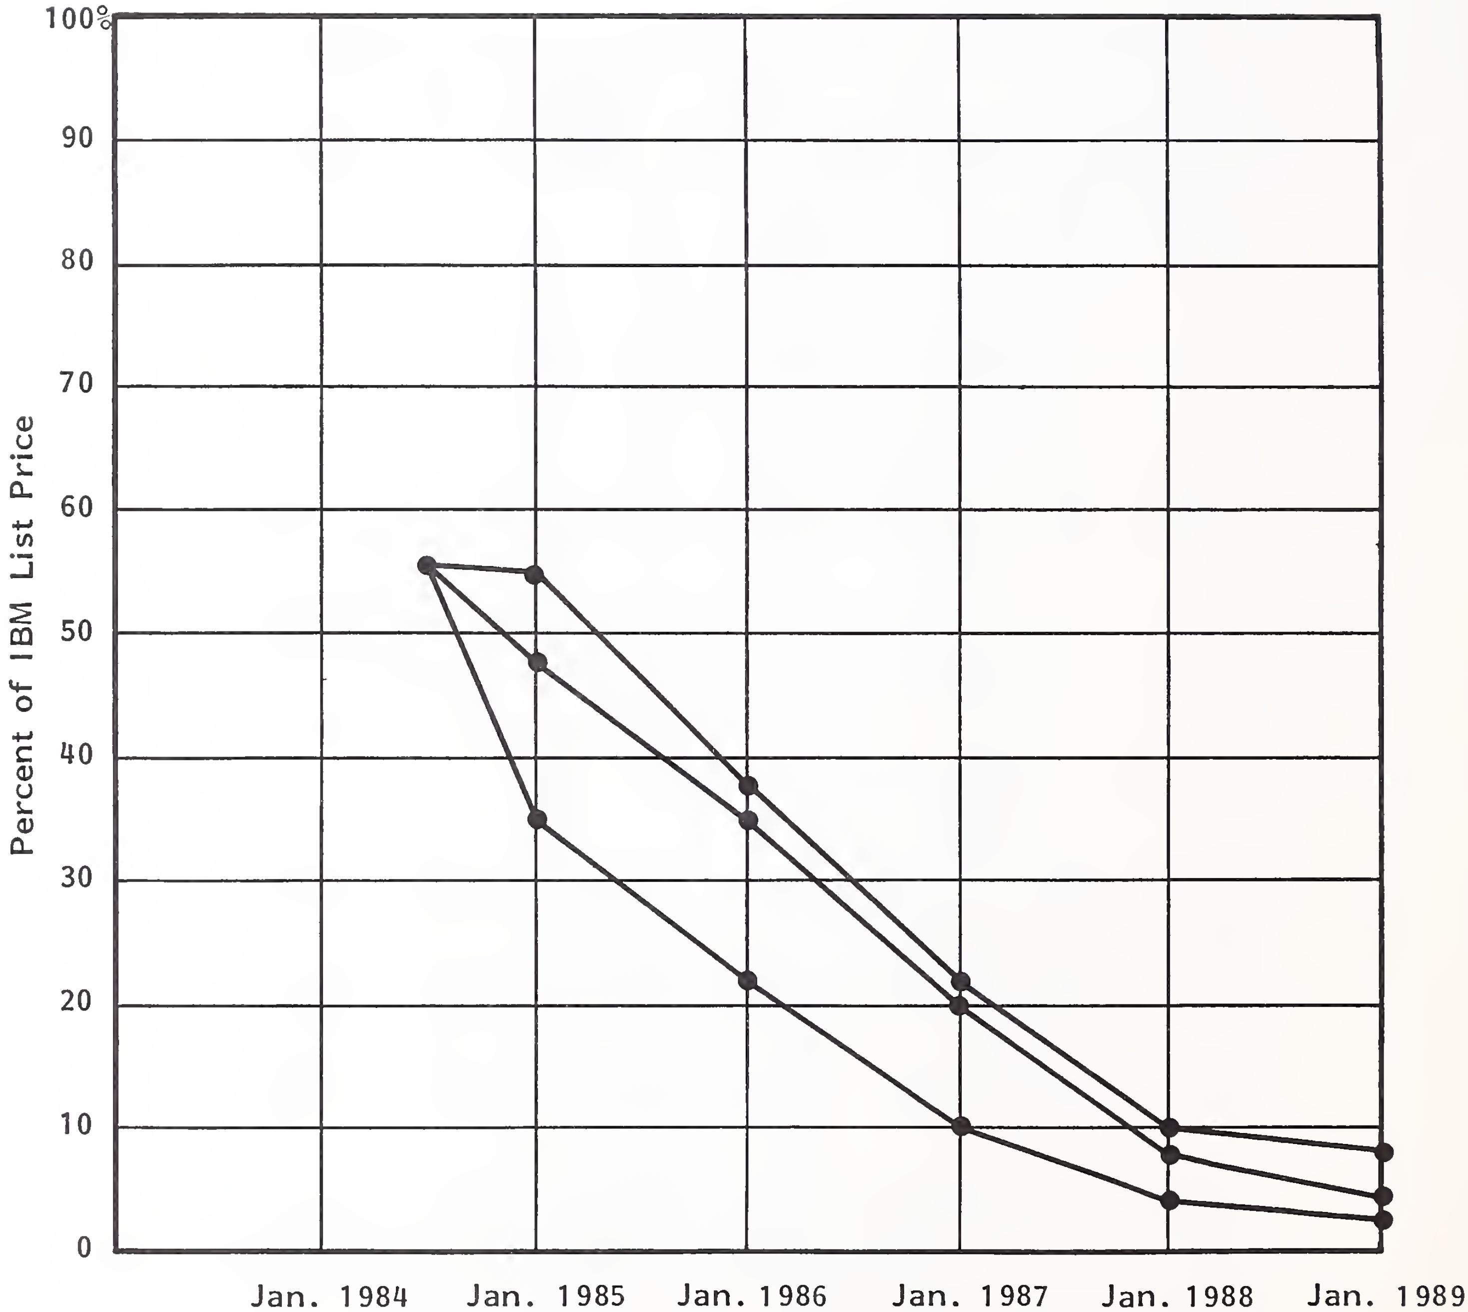

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	54	38	22	10	8
Expected	48	35	20	8	5
Low	35	22	10	5	3

EXHIBIT III-6

RESIDUAL VALUE FORECAST
FOR IBM 4361 PROCESSOR

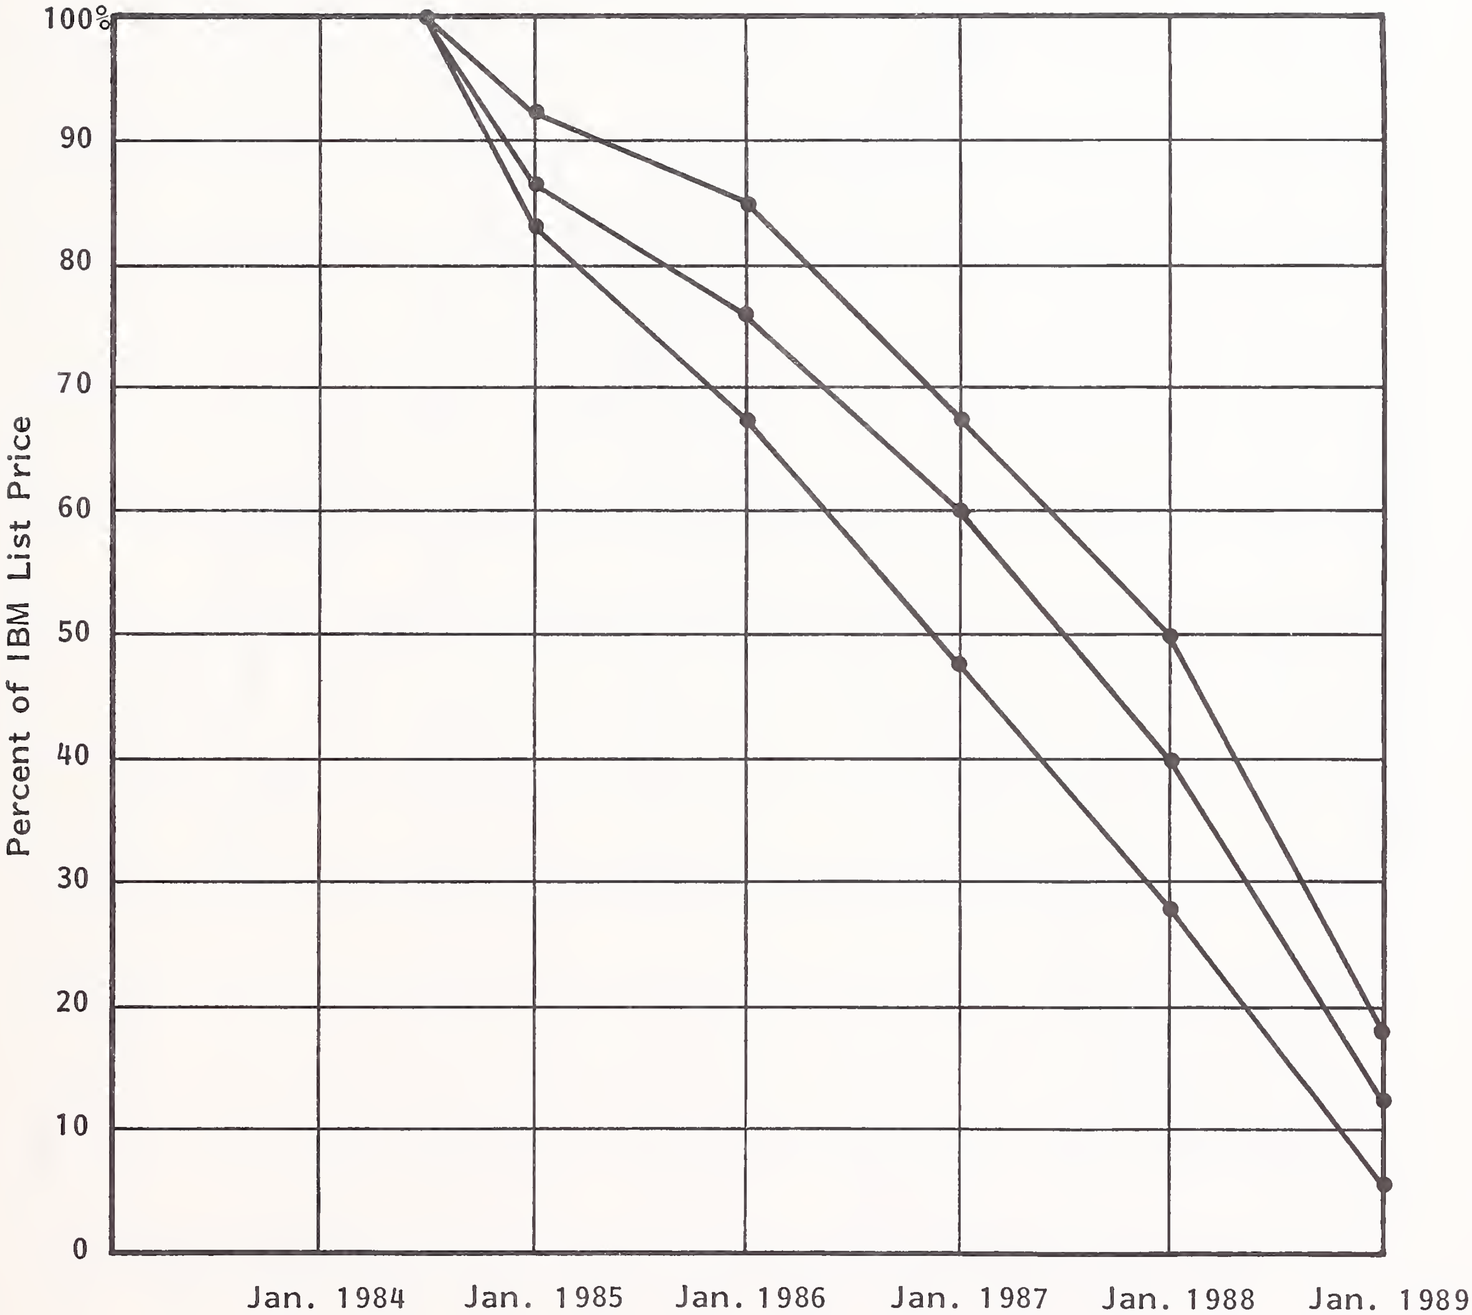

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	92	85	68	50	18
Expected	87	76	60	40	13
Low	84	68	48	28	6

EXHIBIT III-7

RESIDUAL VALUE FORECAST FOR
IBM 4381 PROCESSOR

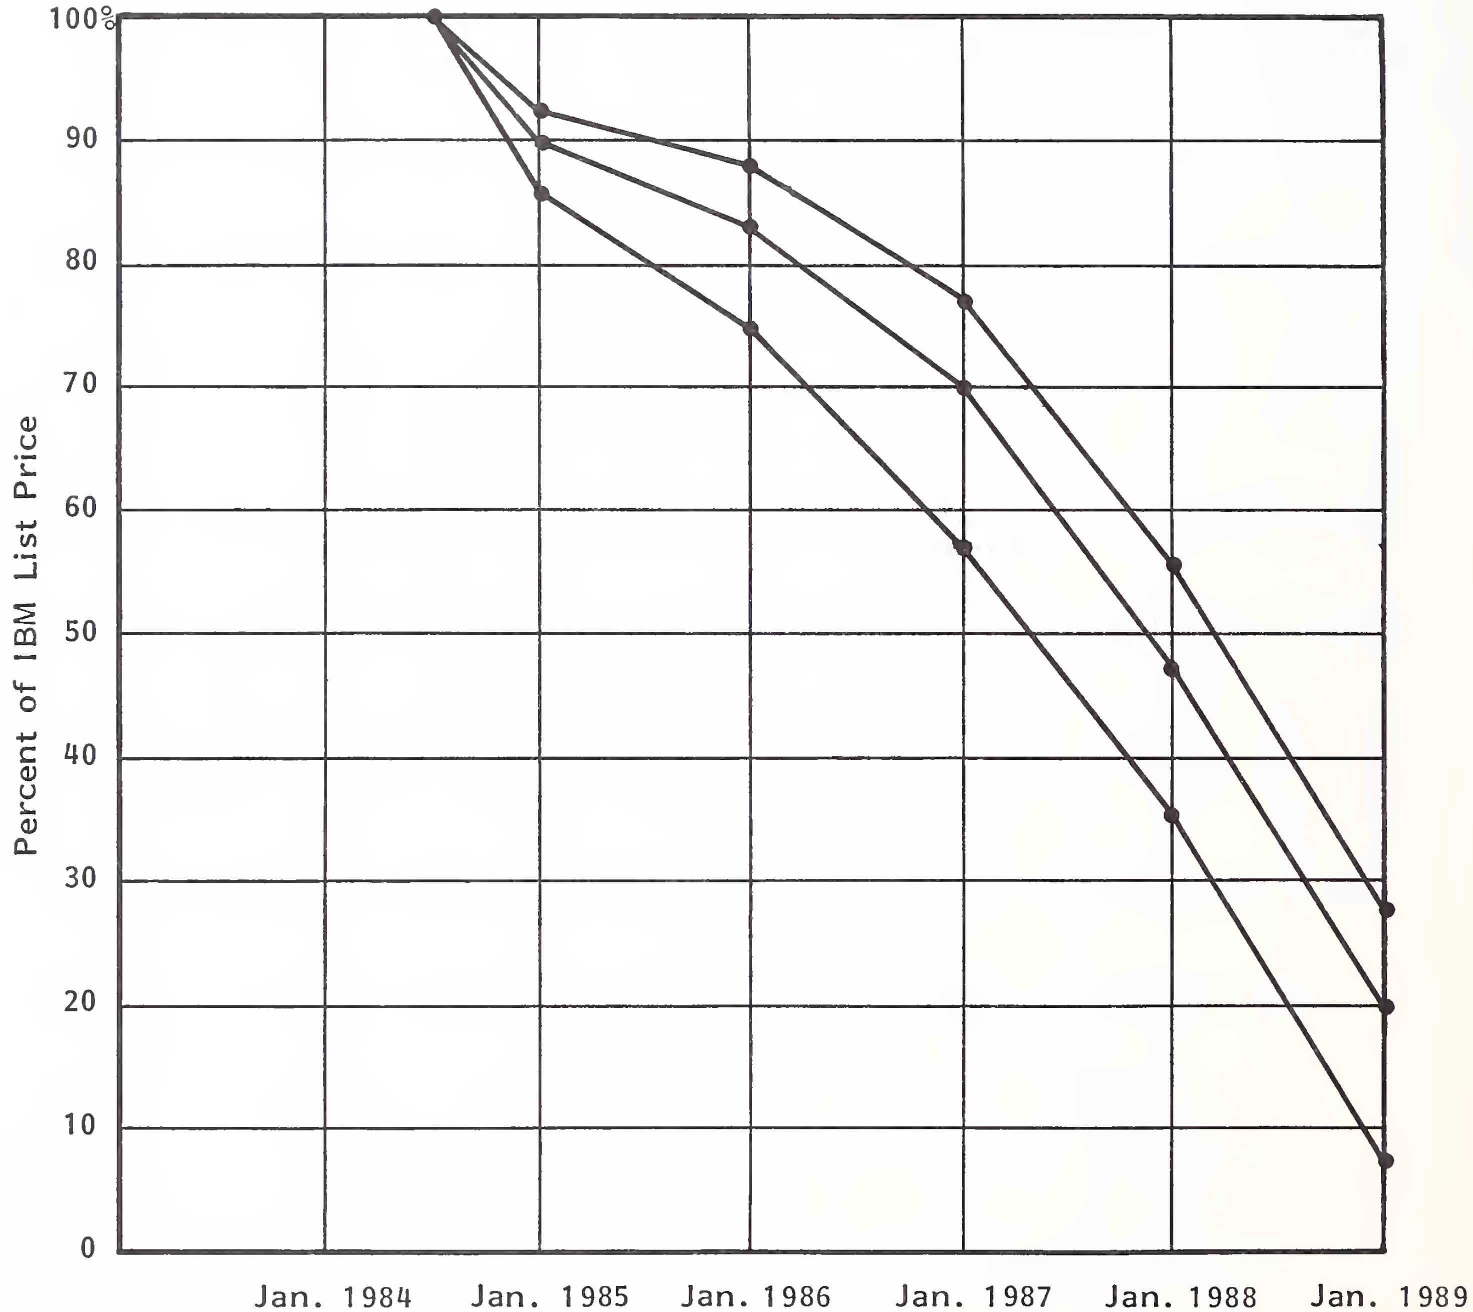

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	92	88	78	56	28
Expected	90	83	70	48	20
Low	86	75	57	36	8

EXHIBIT III-8

RESIDUAL VALUE FORECAST FOR
IBM 3083EX PROCESSOR

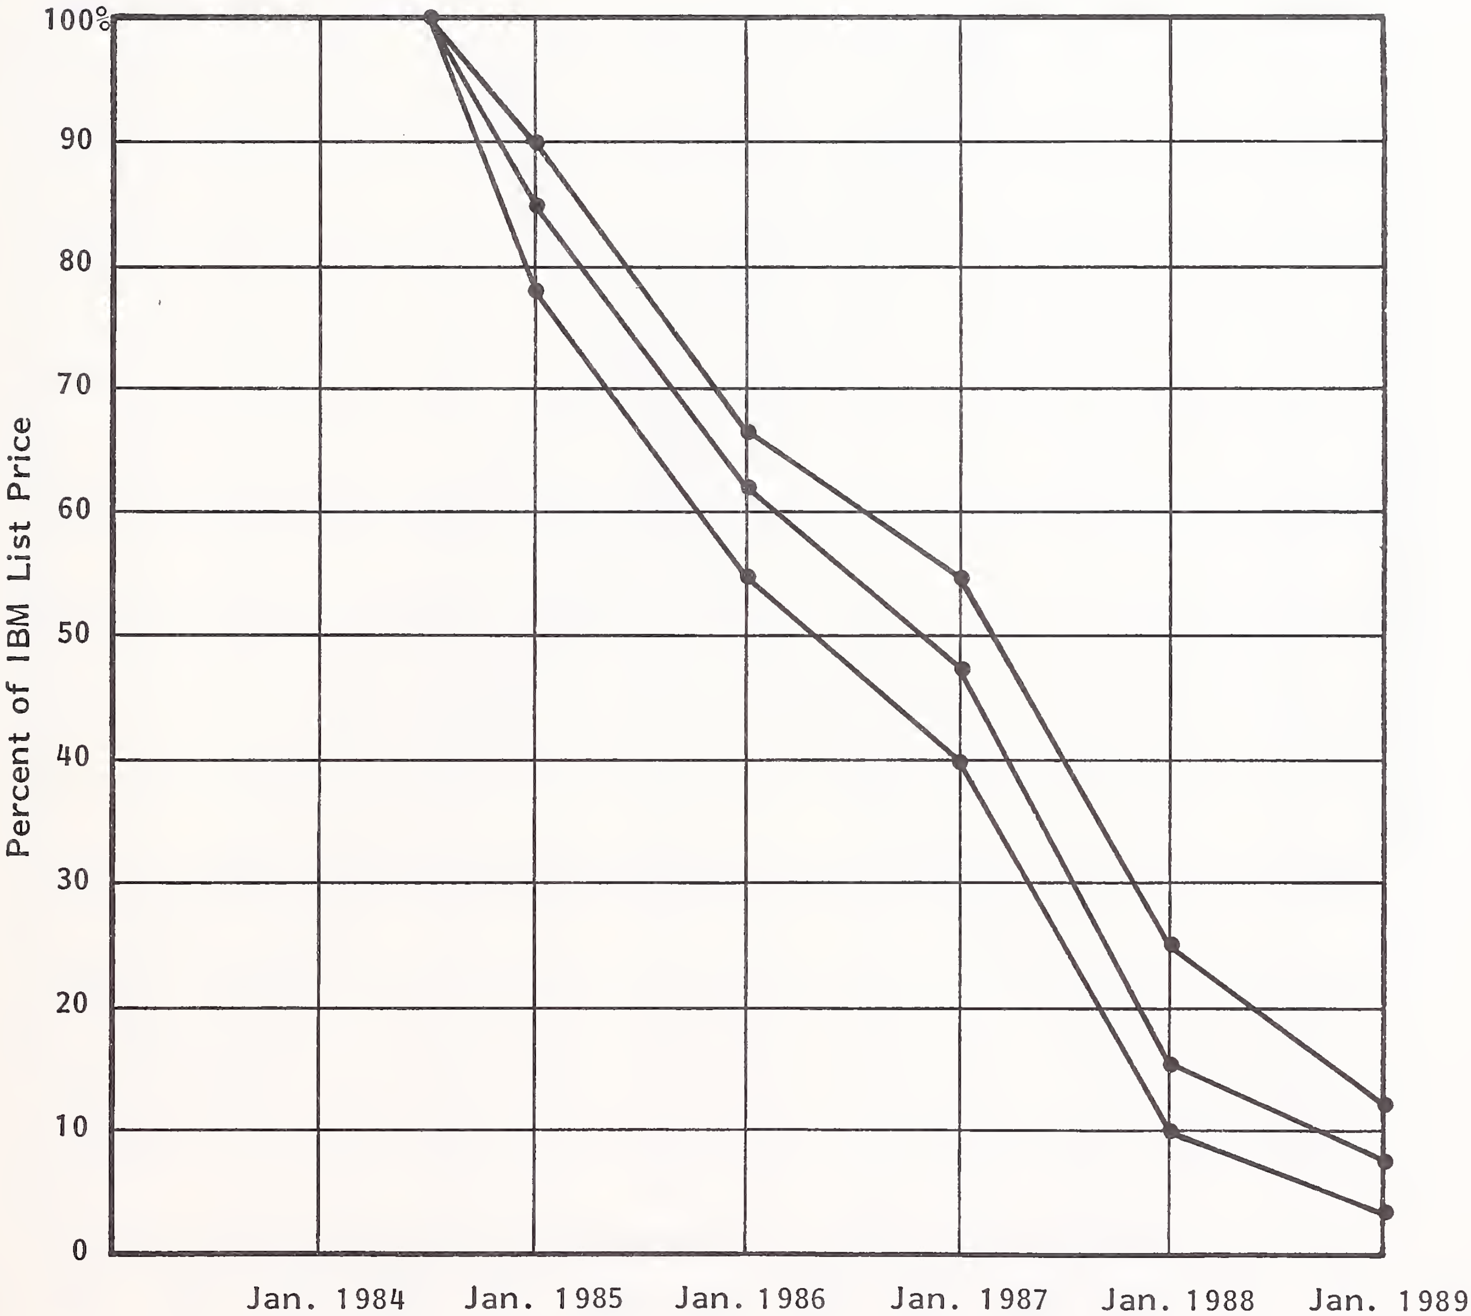

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	90	67	55	25	12
Expected	85	62	48	16	8
Low	78	55	40	10	4

EXHIBIT III-9

RESIDUAL VALUE FORECAST FOR
IBM 3083JX PROCESSOR

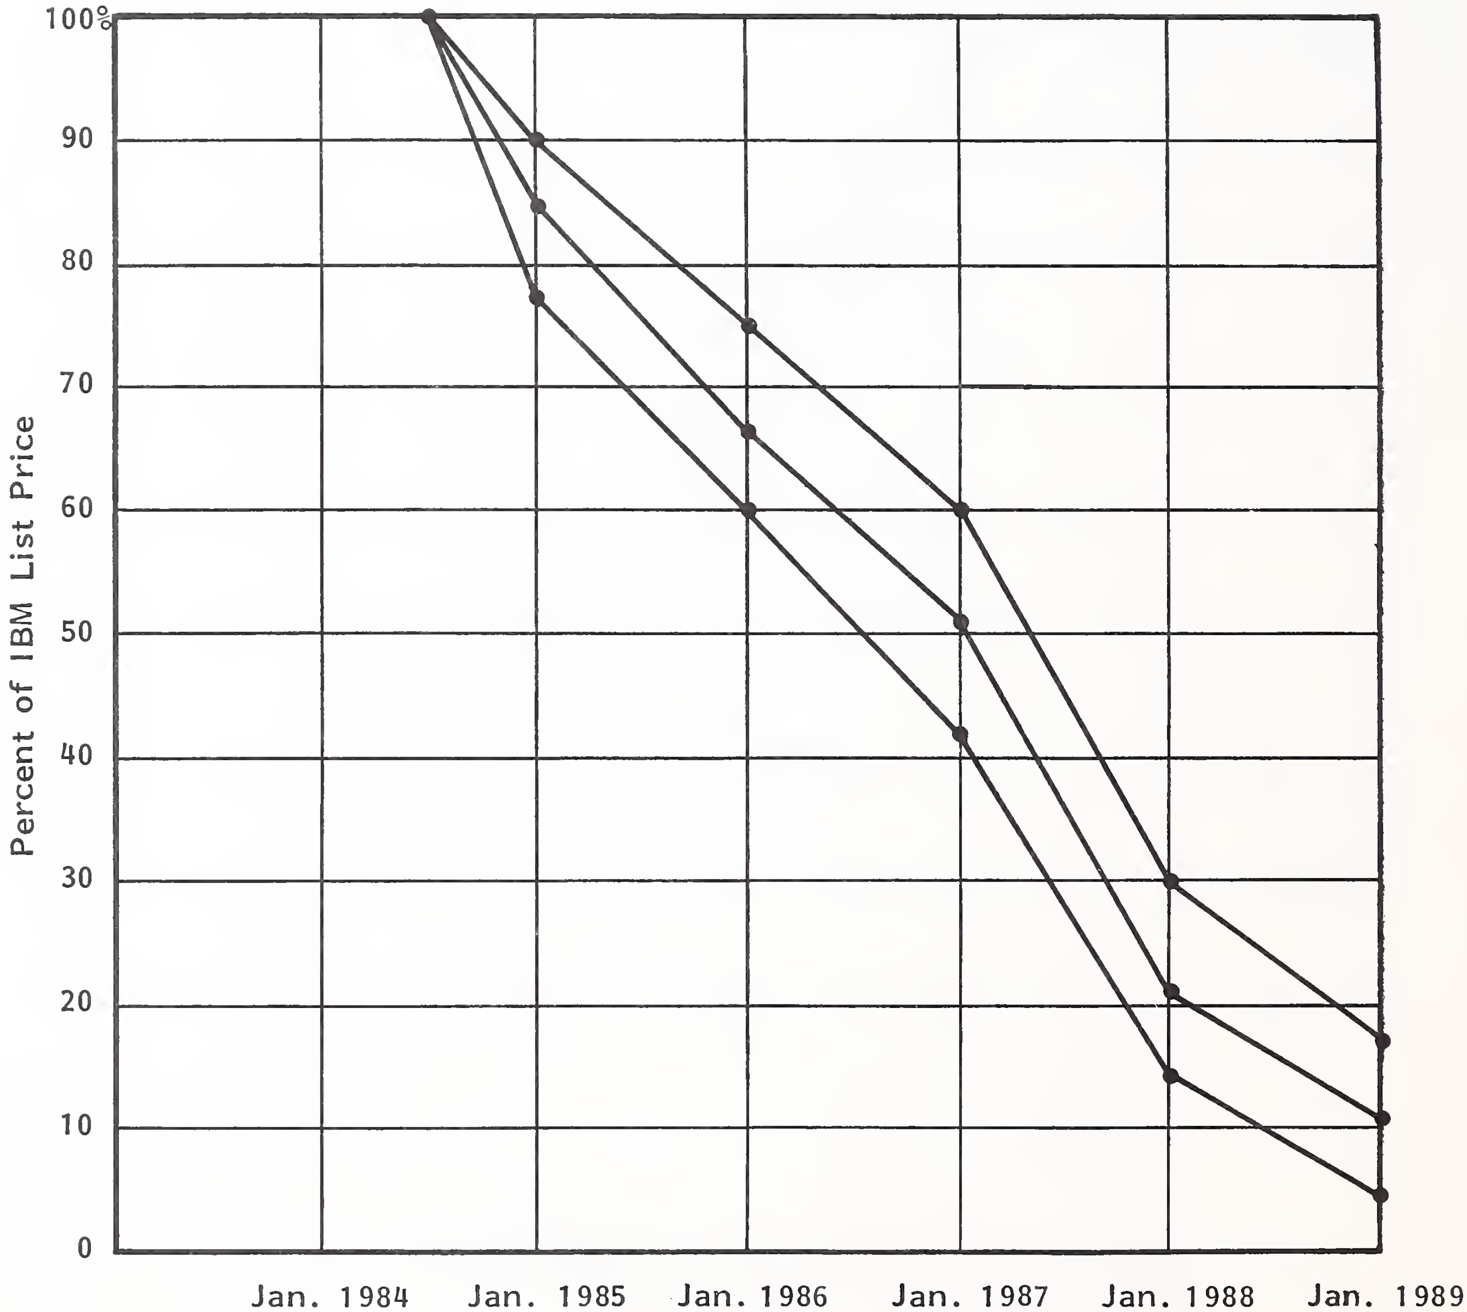

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	90	75	60	30	18
Expected	85	67	51	21	11
Low	78	60	42	15	5

EXHIBIT III-10

RESIDUAL VALUE FORECAST FOR
IBM 3081GX PROCESSOR

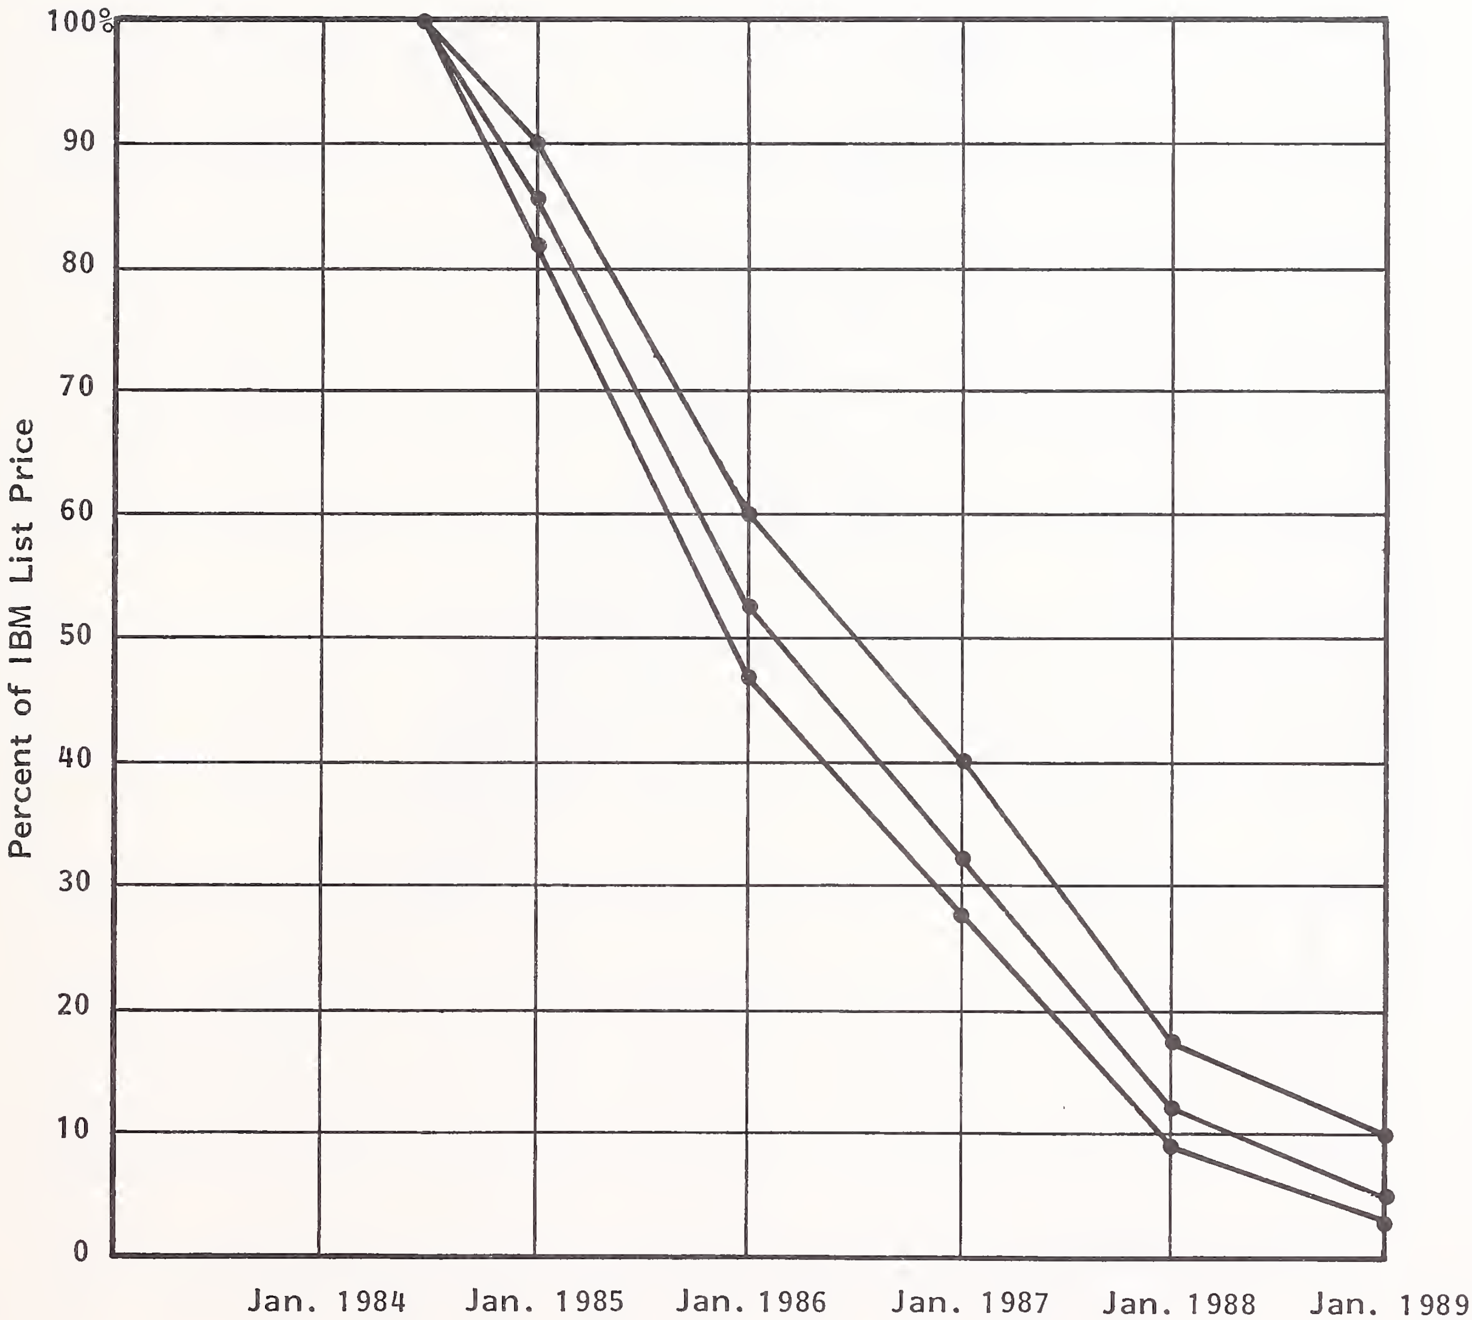

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	90	60	40	18	10
Expected	86	53	32	12	5
Low	82	48	28	9	3

EXHIBIT III-11

RESIDUAL VALUE FORECAST FOR
IBM 3081KX PROCESSOR

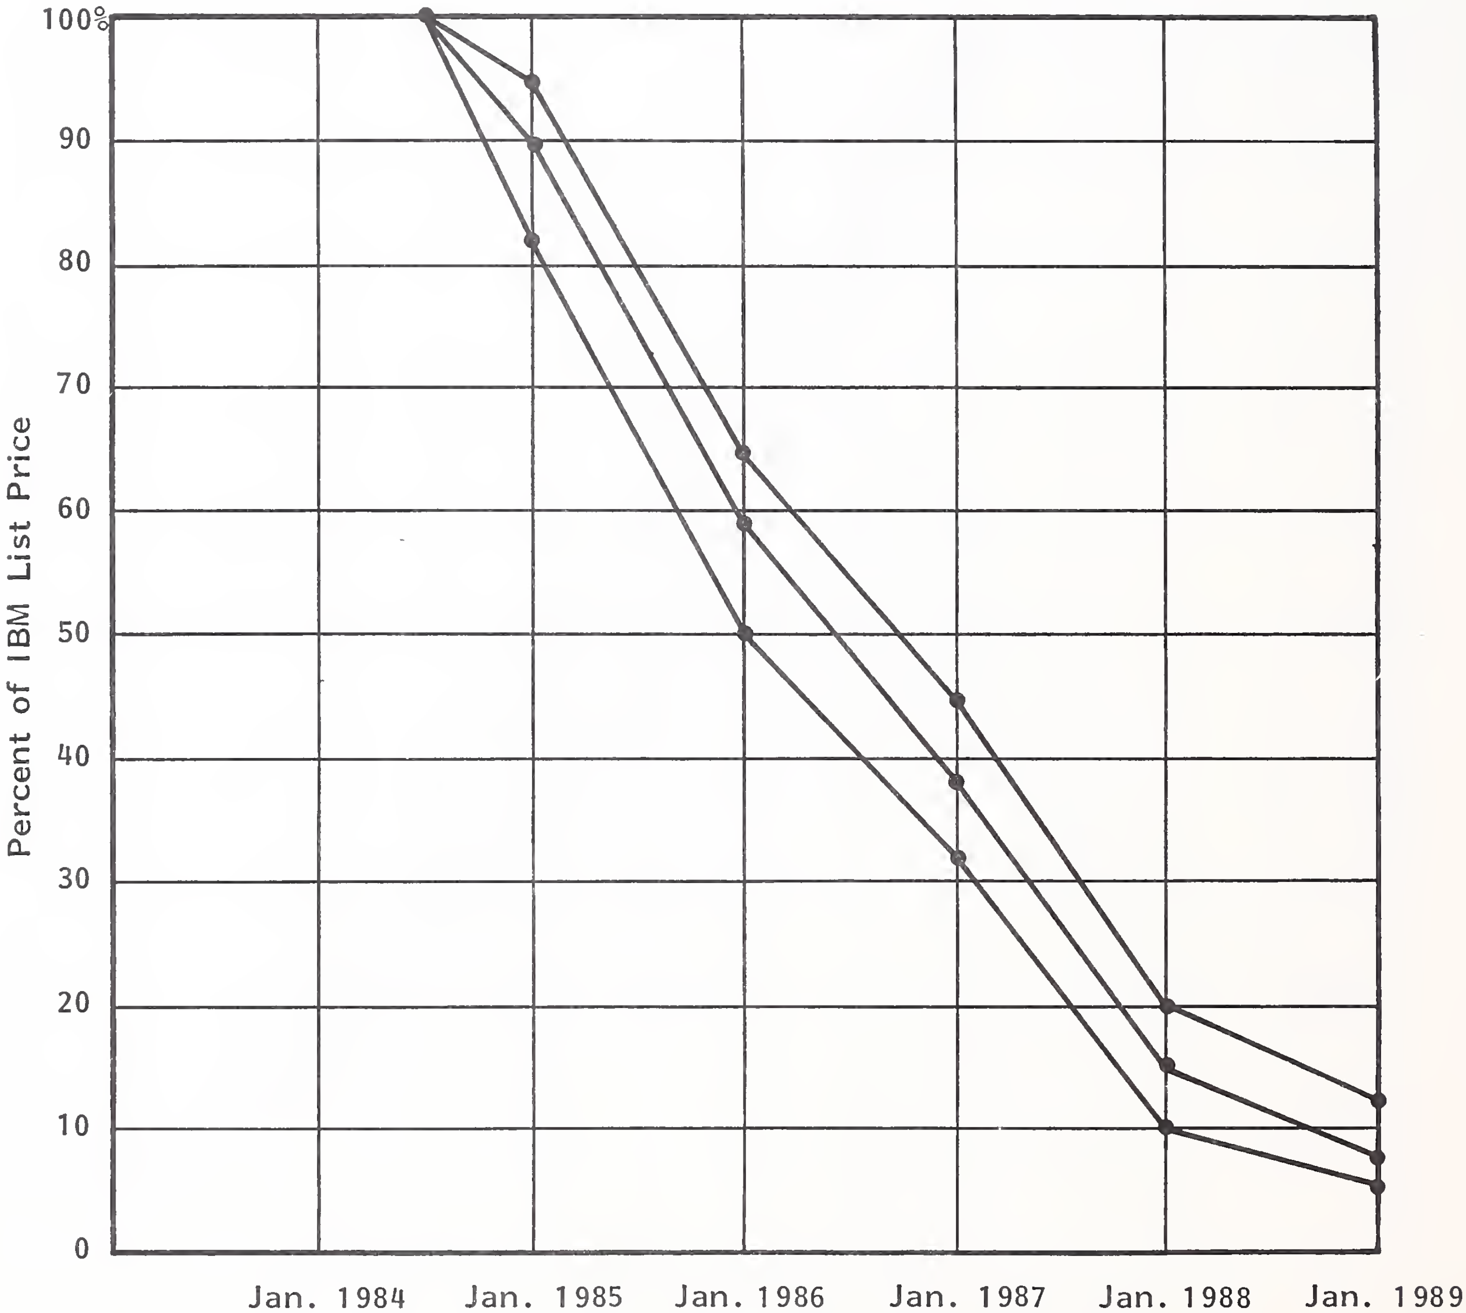

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	95	65	45	20	12
Expected	90	59	38	16	8
Low	82	50	32	10	6

EXHIBIT III-12

RESIDUAL VALUE FORECAST FOR
IBM 3084QX PROCESSOR

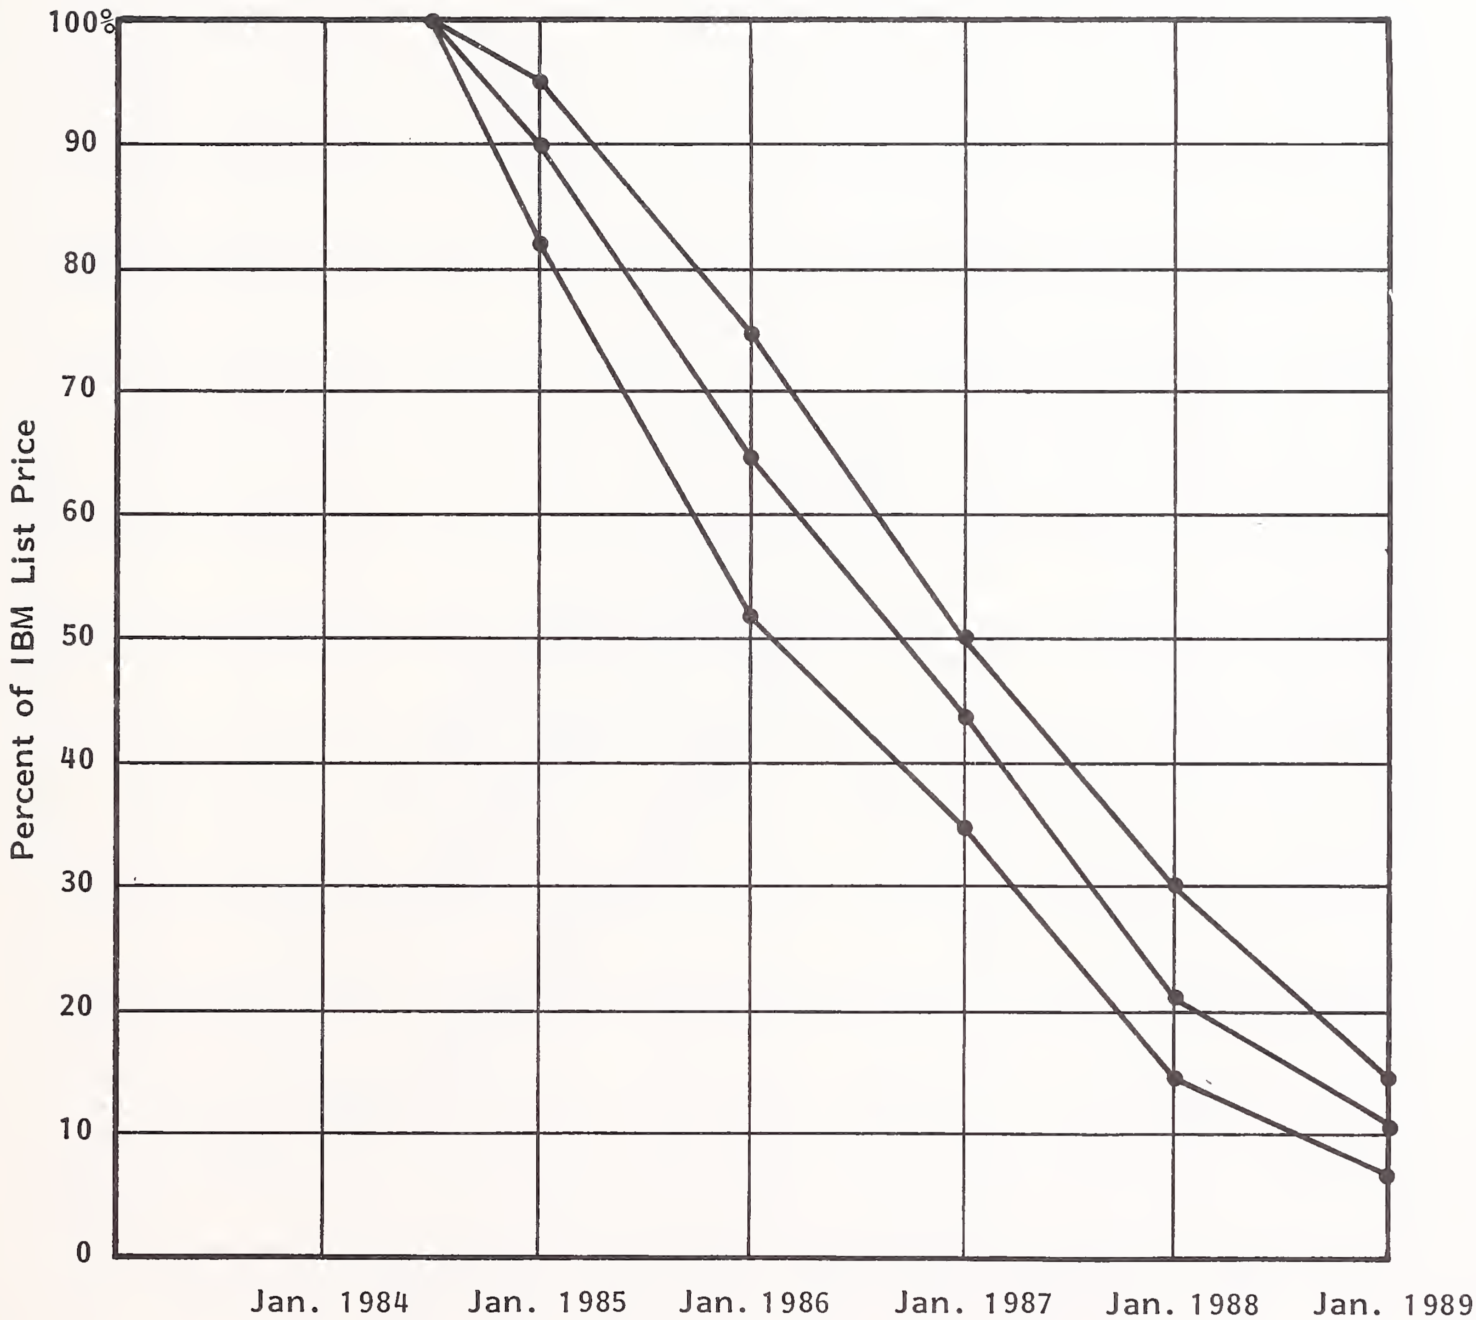

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	95	75	50	30	15
Expected	90	65	44	21	11
Low	82	52	35	15	7

EXHIBIT III-13

RESIDUAL VALUE FORECAST FOR
AMDAHL 5860 PROCESSOR

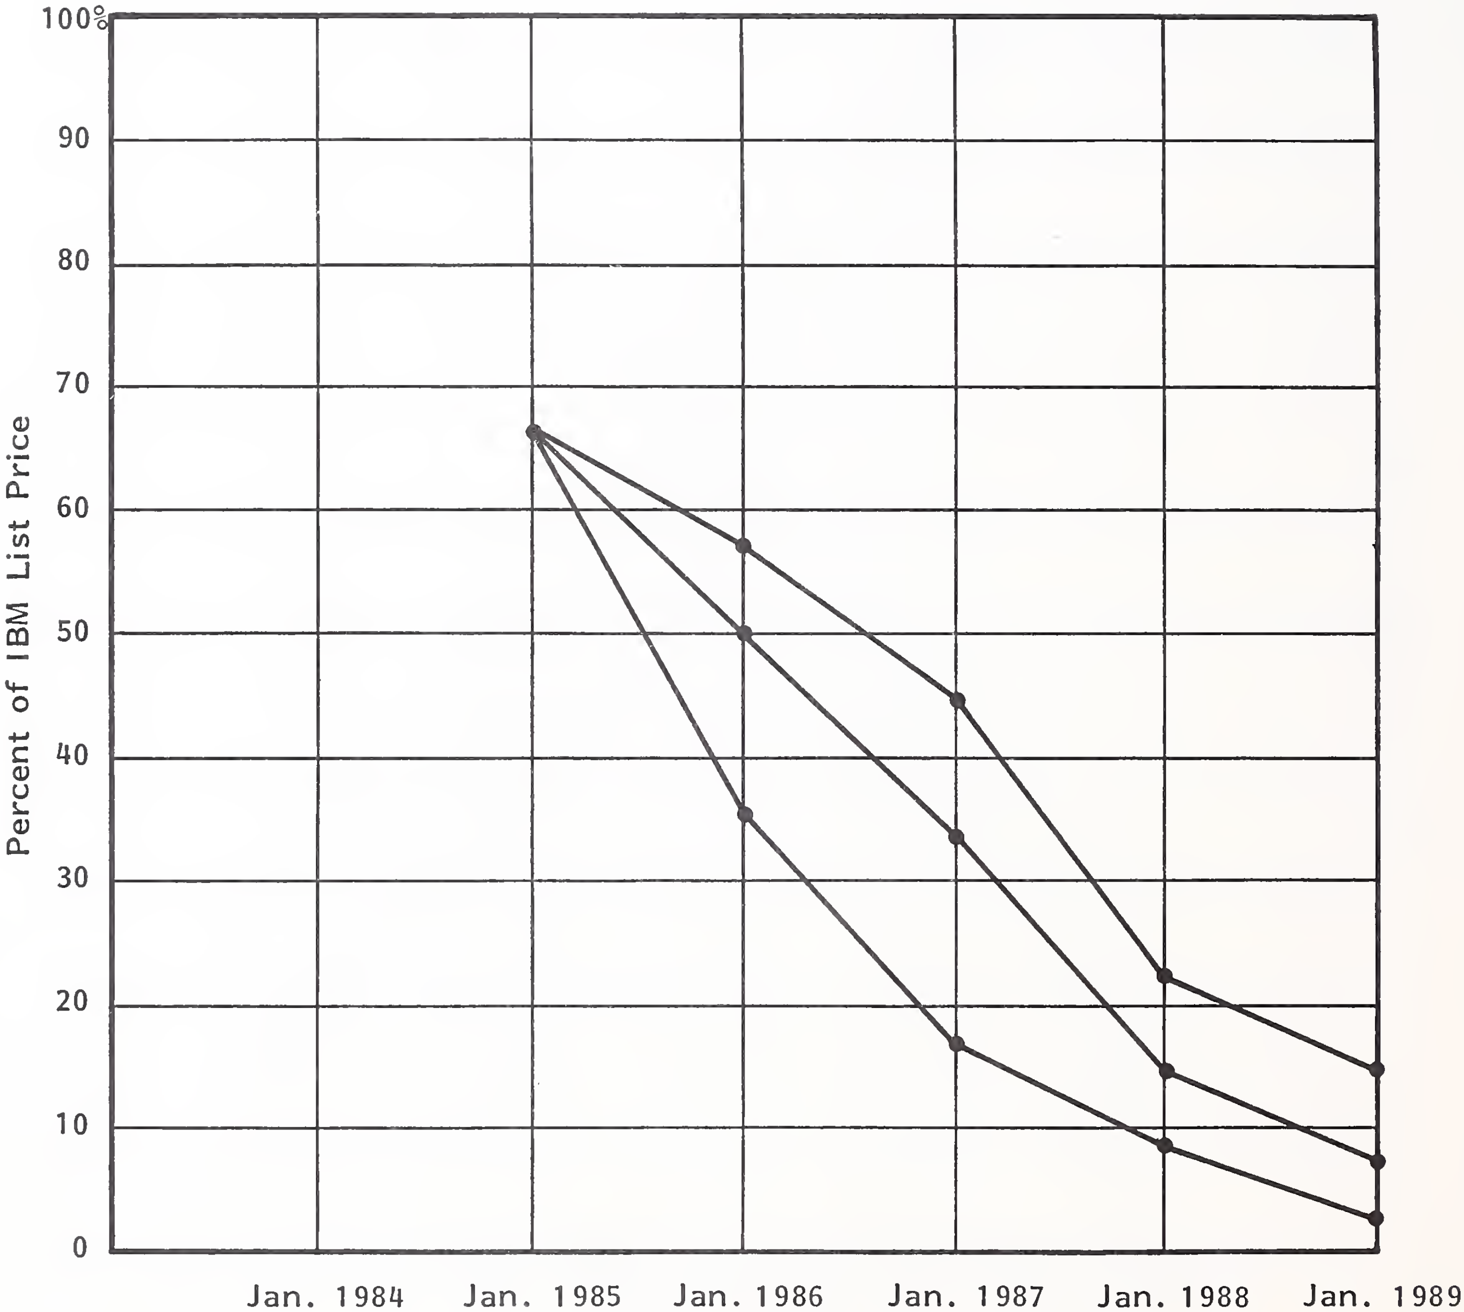

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	–	58	45	22	15
Expected	67	50	34	15	8
Low	–	36	18	9	3

EXHIBIT III-14

RESIDUAL VALUE FORECAST FOR
AMDAHL 5870 PROCESSOR

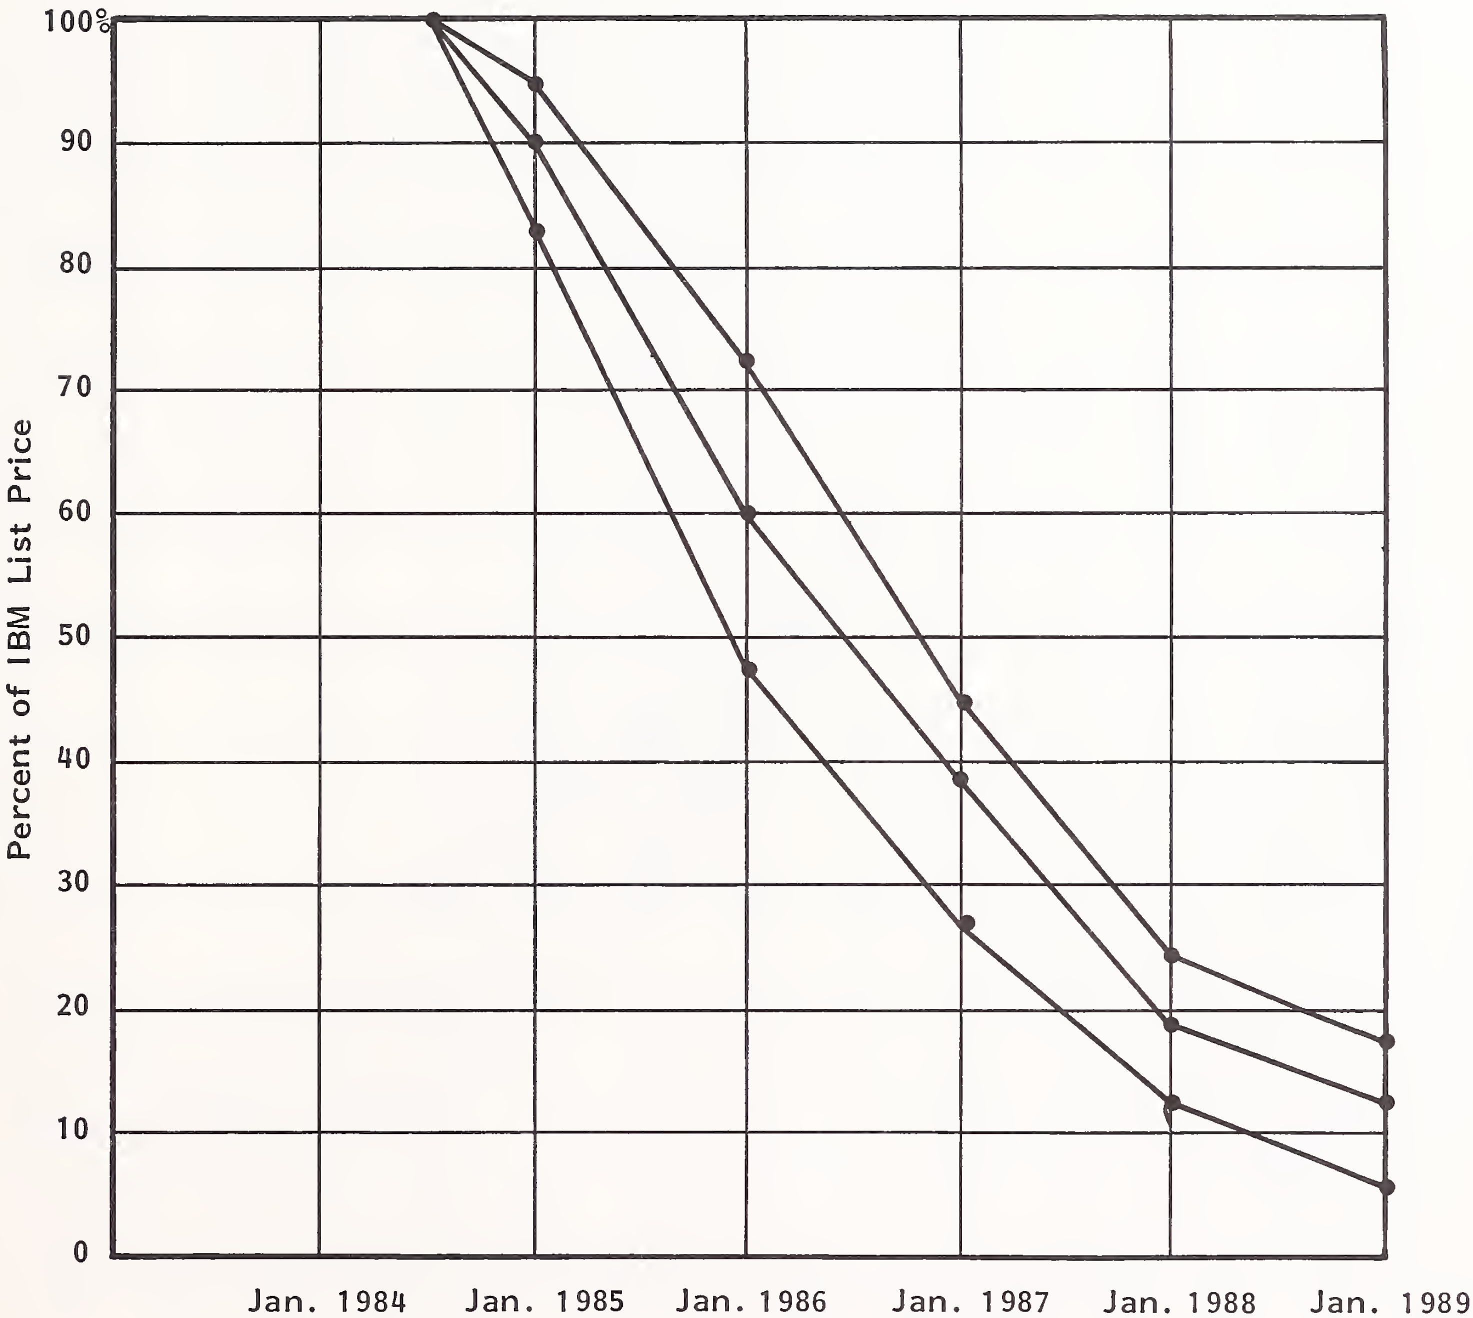

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	95	72	45	25	18
Expected	90	60	39	19	12
Low	83	48	27	12	6

EXHIBIT III-15

RESIDUAL VALUE FORECAST FOR
NAS 6000 SERIES PROCESSOR

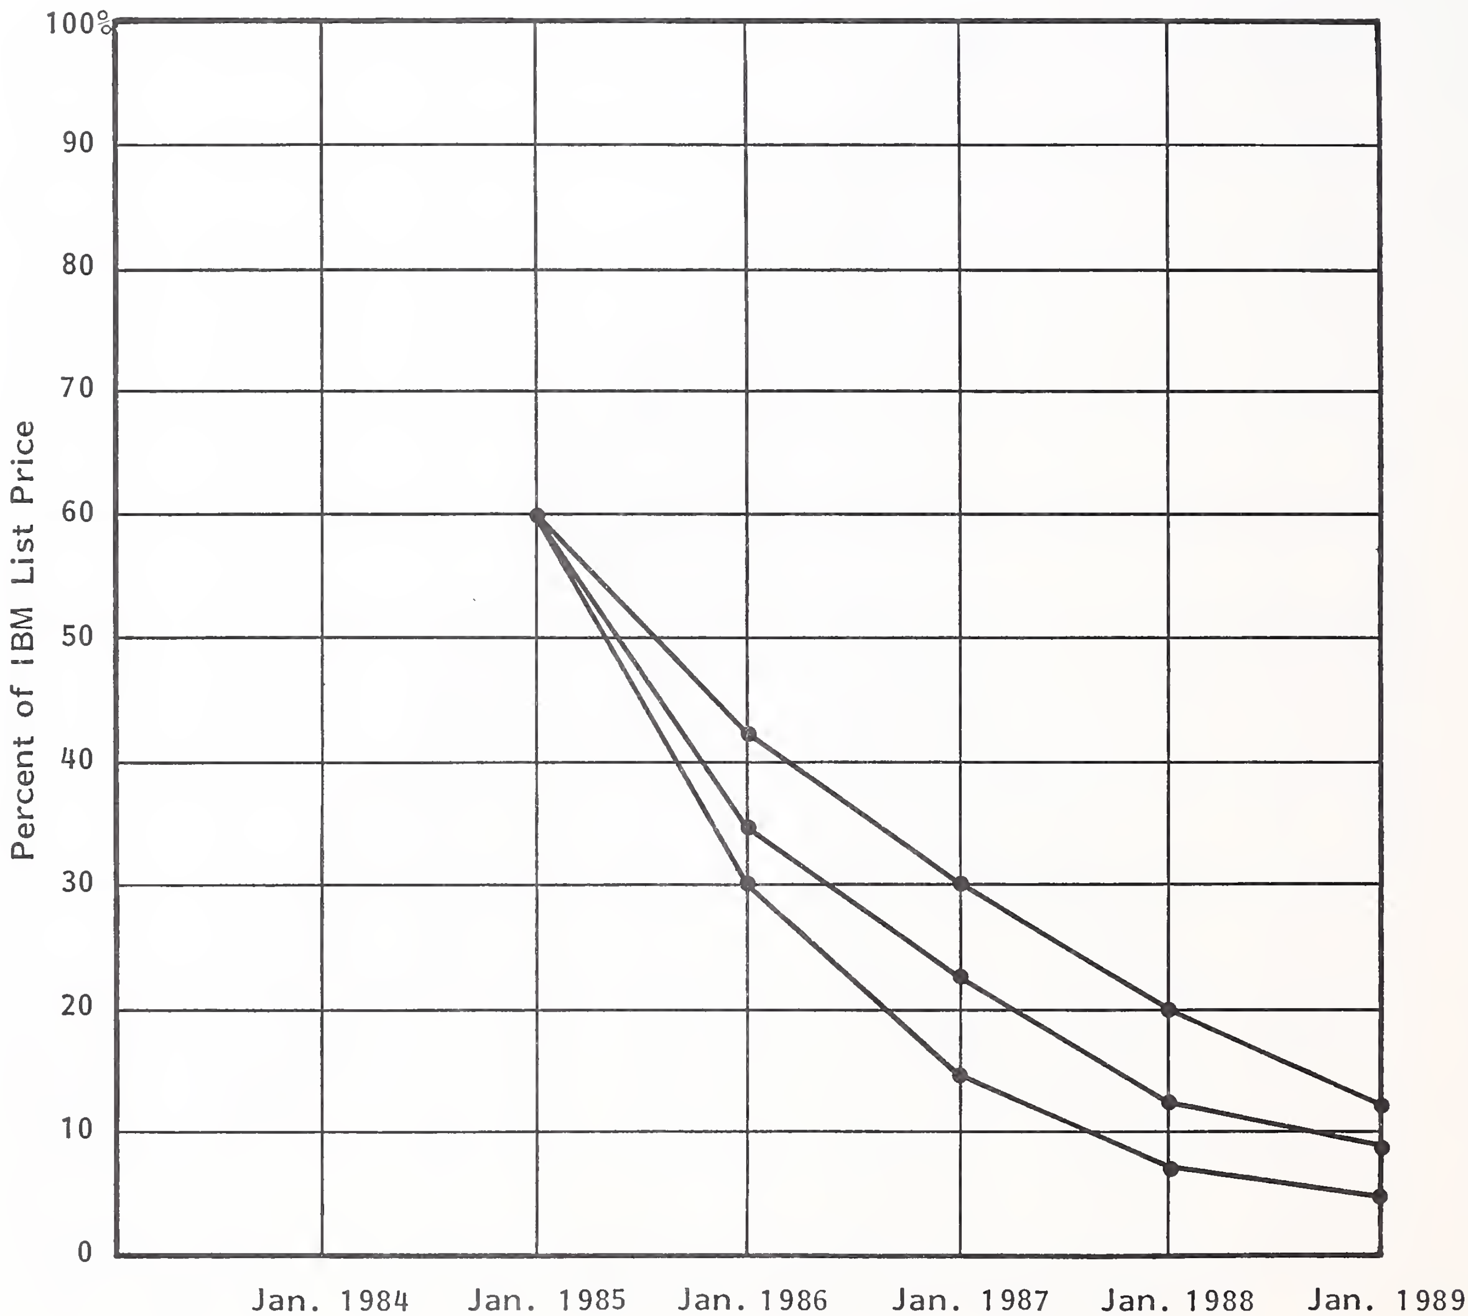

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	–	42	30	20	12
Expected	60	35	22	12	9
Low	–	30	15	8	5

EXHIBIT III-16

RESIDUAL VALUE FORECAST FOR
NAS 8000 SERIES PROCESSOR

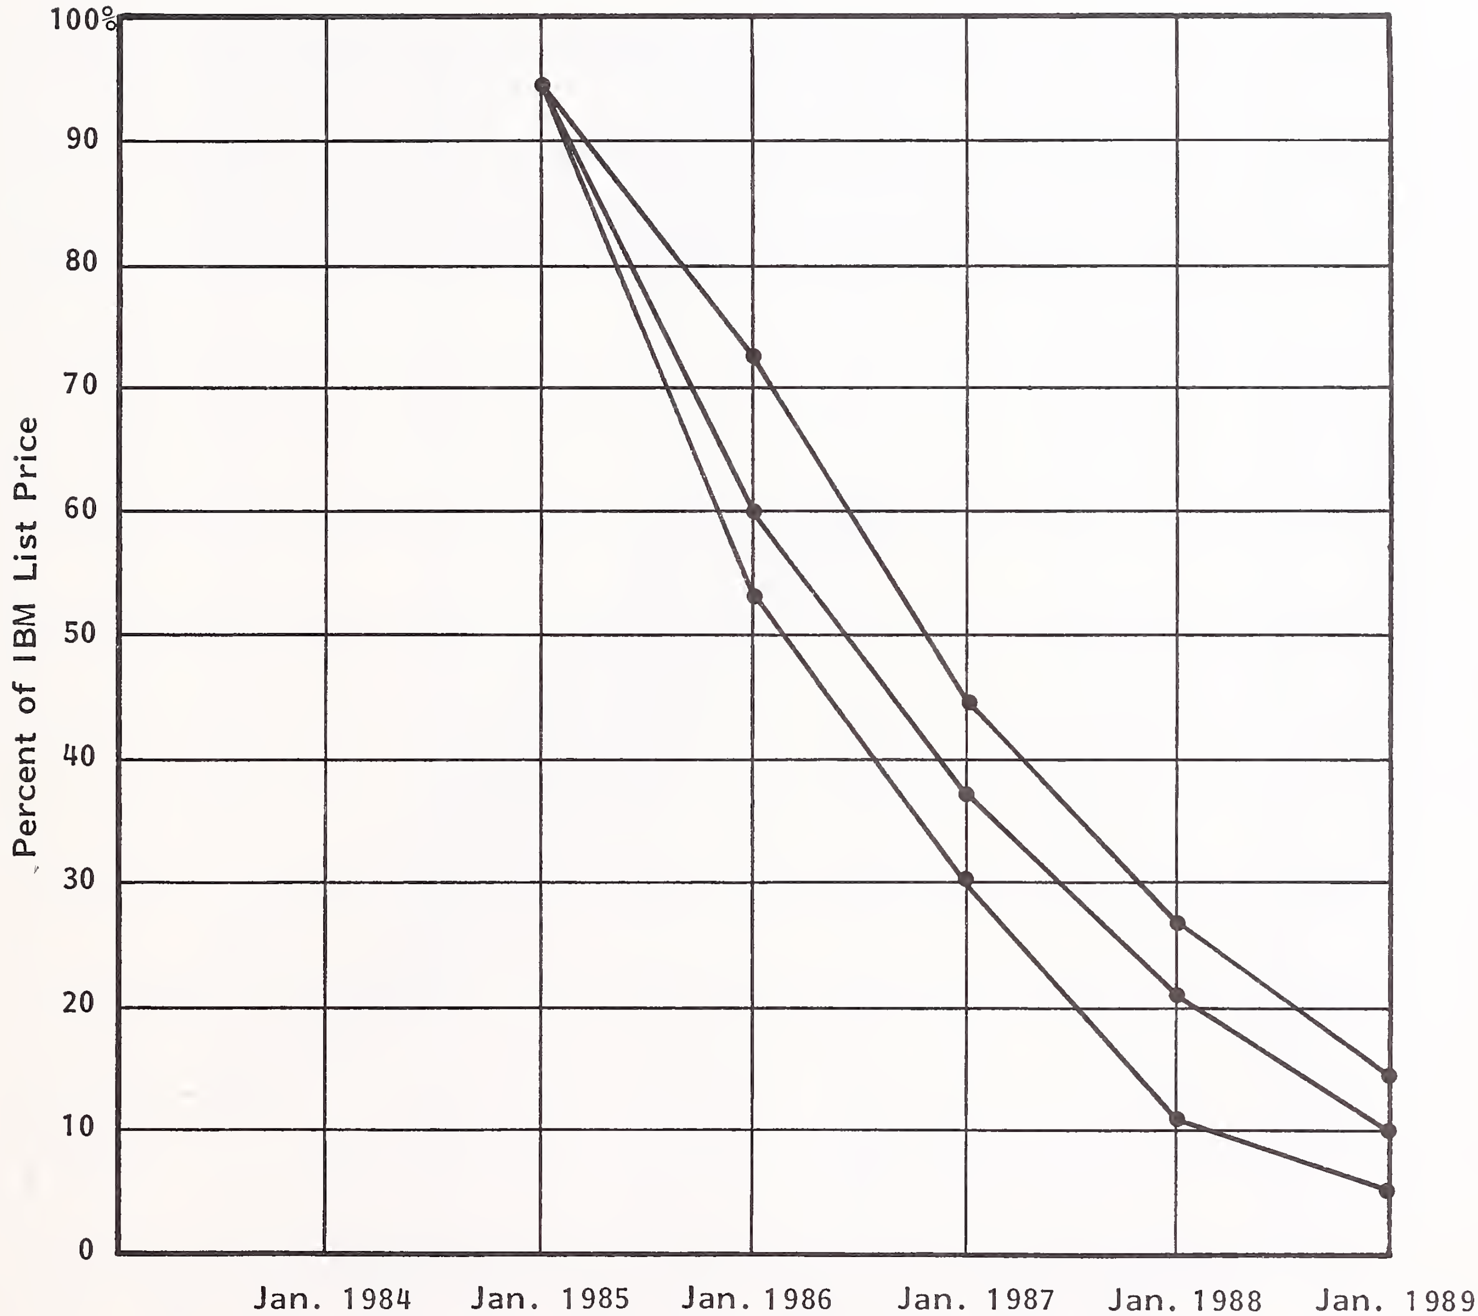

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	–	72	45	28	15
Expected	95	60	38	21	10
Low	–	54	30	11	6

EXHIBIT III-17

RESIDUAL VALUE FORECAST FOR
NAS 9000 SERIES PROCESSOR

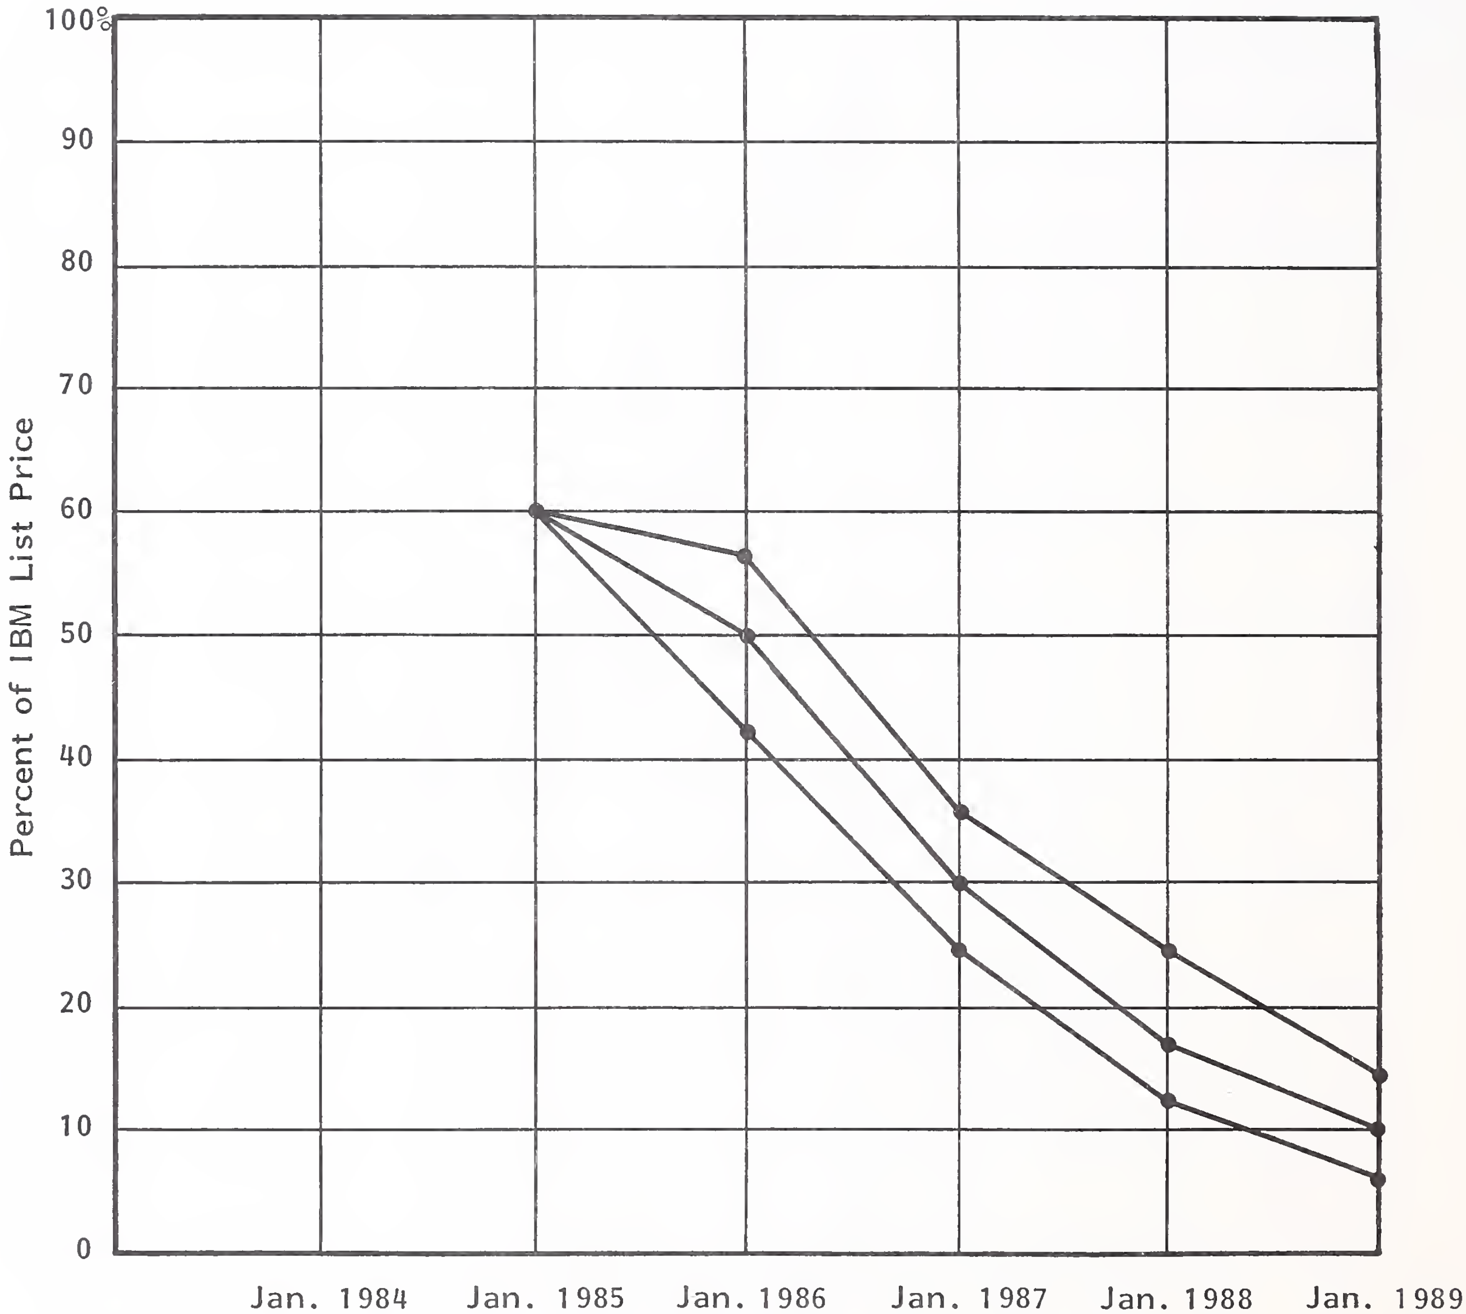

PROJECTED VALUES RANGE	JAN. 1985	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989
High	–	57	36	25	15
Expected	60	50	30	18	10
Low	–	42	25	12	7

- Amdahl 5860 and 5870.

- NAS 6000, 8000, and 9000.

* The values shown are wholesale prices--the amount a used-computer dealer will pay for equipment for subsequent resale to an end user. The factors affecting computer equipment residual values were presented in Residual Value Forecasts for Large-Scale Systems, December 1983. Those factors have received detailed analysis in the past as part of INPUT's residual value series.

INPUT® 1943 Landings Drive, Mountain View, CA 94043, (415) 960-3990